IRY, VOL. 39, NO. 3, MAY 19 8097 IEEE Transactions on Information Theory 39(1993)May, No.3, New York, US

Optimal Nonuniform Signaling for Gaussian Channels

Frank R. Kschischang, *Member, IEEE,* and Subbarayan Pasupathy, *Fellow, IEEE*

Abstract— **Variable-rate data transmission schemes in which constellation points are selected according to a nonuniform probability distribution are studied. When the criterion is one of minimizing the average transmitted energy for a given average bit rate, the best possible distribution with which to select constellation points is a Maxwell–Boltzmann distribution. In principle, when constellation points are selected according to a Maxwell–Boltzmann distribution, the ultimate shaping gain ($\pi e/6$ or 1.53 dB) can be achieved in *any* dimension. Nonuniform signaling schemes can be designed by mapping simple variable-length prefix codes onto the constellation. Using the Huffman procedure, prefix codes can be designed that approach the optimal performance. These schemes provide a fixed-rate primary channel and a variable-rate secondary channel, and are easily incorporated into standard lattice-type coded modulation schemes.**

Index Terms—**Signal constellations, maximum entropy principle, shaping gain, coded modulation, Huffman codes.**

I. Introduction

In the conventional approach to data transmission, each point in a given constellation is equally likely to be transmitted. While this approach yields the maximum bit rate for a given constellation size, it does not take into account the energy cost of the various constellation points. In this paper, the idea of choosing constellation points with a nonuniform probability distribution is explored. Such nonuniform signaling will reduce the entropy of the transmitter output, and hence the average bit rate. However, if points with small energy are chosen more often than points with large energy, energy savings may (more than) compensate for this loss in bit rate.

It follows immediately from the maximum entropy principle (see, e.g., [1, ch. 11]), or by variational calculus as in [2, Section IV-B], that the probability distribution that maximizes entropy for a fixed average energy is one in which a constellation point $\boldsymbol{r}$, with energy $\|\boldsymbol{r}\|^2$, is chosen with probability $p(\boldsymbol{r}) \propto \exp(-\lambda\|\boldsymbol{r}\|^2)$, where the nonnegative parameter λ governs the trade-off between bit rate and average energy. Nonuniform signaling with this family of distributions, well known in statistical mechanics and thermodynamics as Maxwell–Boltzmann distributions [3], [4] or as Gibbs ensembles [5], is the focus of this paper.

Nonuniform signaling is closely related to the notion of constellation shaping in coded modulation (as described in, e.g., [2], [6]–[9]). Constellation shaping can provide an energy savings called shaping gain in addition to the usual coding gain provided by lattice- or trellis-coding. Indeed, the gain G provided by a coded modulation system (relative to a simple pulse amplitude modulation baseline) operating at a bit rate β (bits per two-dimensional channel-use) is well-approximated by writing

$$G \approx \gamma_c\gamma_s(1-2^{-\beta}),$$

where γ_c and γ_s denote, respectively, the coding gain [10] and shaping gain [2] of the scheme in question. As will be shown, the discretization factor $1-2^{-\beta}$ properly adjusts the gain for finite bit rates. The connection between shaping and nonuniform signaling arises from the fact that, in schemes that employ shaping, a nonuniform distribution is induced on the points of the low-dimensional constituent constellation. By applying nonuniform signaling directly (rather than indirectly via constellation shaping), nonuniform signaling can, in any dimension, achieve the ultimate shaping gain—$\pi e/6$ or 1.53 dB—attainable with uniform signaling only in the limit of infinite dimension [2]. Indeed, one of the principal results of this paper is a method of designing simple nonuniform signaling schemes that approach this ultimate level.

It is important to note at the outset, however, that practical implementation of direct nonuniform signaling will be hampered by the variable transmission rate of such schemes. Transmitting data obtained from a fixed-rate source requires data buffering at the transmitter and the receiver, which leads to the problem of coping with buffer over- or underflow. Furthermore, since the transmitted signals represent variable numbers of bits, channel errors may cause the insertion and deletion of bits in the decoded data, causing potential losses of synchronization. Although these system problems will tend to limit the broad applicability of nonuniform signaling, we do not attempt to provide solutions to these problems in this paper.

Instead, our aim in this paper is 1) to provide insight into nonuniform signaling schemes and how they relate to conventional signaling schemes, 2) to assess the potential gains that nonuniform signaling may provide, and 3) to provide a method (via the Huffman algorithm) by which simple, near-optimal, nonuniform signaling schemes may be designed. Such nonuniform signaling schemes provide a standard, fixed-

Manuscript received December 28, 1990; revised August 20, 1992. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada, and by the Government of Ontario through an Ontario Graduate Scholarship. This paper was presented in part at the 1991 Tirrenia International Workshop on Digital Communications, Tirrenia, Italy, September 8–12, 1991, and at the 16th Biennial Symposium on Communications, Kingston, ON, Canada, May 27–29, 1992.

The authors are with the Department of Electrical Engineering, University of Toronto, Toronto, ON, Canada M5S 1A4.

IEEE Log Number 9205742.

rate, primary channel unaffected by the system problems mentioned in the previous paragraph, together with a variable-rate secondary channel, in which these system problems may be acceptable. At various places in this paper, the close analogy between these nonuniform signaling schemes and mathematically equivalent statistical mechanical systems will be pointed out.

To motivate our interpretation of nonuniform signaling as a shaping method, we summarize the complementary notions of coding and shaping in Section II. In Section III, we briefly define the coded modulation parameters that will be needed throughout the paper. In Section IV, we discuss various properties of the Maxwell–Boltzmann distribution that are important in this setting. The results of applying this distribution to signal point selection from spherical constellations based on various dense lattices are given in Section V. In Section VI, we define and apply "continuous approximations" to show that the ultimate shaping gain of $\pi e/6$, or 1.53 dB, is attainable in any dimension when constellation points are selected with a Maxwell–Boltzmann distribution. In Section VII, we show how the Huffman procedure may be used to obtain dyadic approximations to the Maxwell–Boltzmann distribution that provide near-ultimate shaping gains. In Section VIII, we discuss the integration of nonuniform signaling with coset coding. Finally, we make some general comments and concluding remarks in Section IX.

II. Coding and Shaping

Coding and shaping are two separate and complementary operations that contribute to the gain of lattice-type coded modulation schemes such as lattice codes and lattice-type trellis codes. In implementation as well as in analysis, the two operations are dual and separable and provide two additive gain components: coding gain and shaping gain. We say that coding gain is a *distance property* of the coded modulation scheme because it depends, in general, on the set of distances between the various transmitted signal sequences. Coding is generally performed to achieve a large minimum distance between signal-sequences, i.e., coding attempts to be distance-maximizing. Shaping gain, on the other hand, is an *energy property* of the coded modulation scheme because it depends, in general, on the energy of the various transmitted signal sequences. Shaping is generally performed to achieve a small average transmitted energy while maintaining the desired bit rate, i.e., shaping attempts to be energy-minimizing.

In more general terms, coding and shaping attempt to solve two related, but different, problems. The coding problem is to find a large set of symbol sequences that can be distinguished with high reliability in the presence of noise. The shaping problem is to use these symbol sequences to deliver maximum information to the receiver at minimum cost where, in this paper, cost is measured by the average energy per transmitted symbol.

Most lattice-type coded modulation schemes based on a lattice partition Λ/Λ' have the encoder structure shown in Fig. 1 (see [10]–[12]). The encoder consists of a coset code $\mathbf{C}$ and a signal point selector $\mathbf{S}$. The coset code $\mathbf{C}$ produces

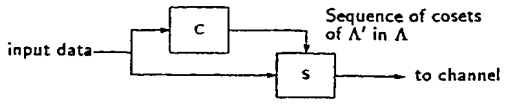

Fig. 1. Encoder structure for lattice-type coded modulation schemes.

a sequence of sets of signals, drawn from the cosets of a sublattice Λ' in a lattice Λ. In general, each coset contains an infinite number of points. The role of the signal point selector $\mathbf{S}$ is to choose an actual constellation point to be transmitted from each infinite coset. The coset code $\mathbf{C}$ is usually chosen to maximize the minimum distance between different possible signal sequences. The signal point selector $\mathbf{S}$ usually attempts to minimize the average transmitted energy while supporting the desired bit rate. In light of our previous discussion, it should be clear that $\mathbf{C}$ performs the coding (or distance-maximizing) operation, while $\mathbf{S}$ performs the shaping (or energy-minimizing) operation. Both the coset code $\mathbf{C}$ and the signal point selector $\mathbf{S}$ contribute to the transmission of data in these coded modulation schemes. In this paper, the signal point selection or shaping component of these coded modulation schemes is studied.

Shaping schemes (or signal point selectors) may be classified as being either fixed-rate or variable-rate. Fixed-rate schemes achieve the transmission of a fixed number of bits over some well-defined signaling interval. Generalized cross constellations [2], Voronoi constellations [6], block shaping codes [7], trellis shaping codes [8], and truncated polydisc constellations [9] are all examples of fixed-rate shaping schemes. From the point of view of maximizing shaping gain, the best possible N-dimensional constellation shape is the N-sphere, since it achieves a specified volume with least average energy.

With variable-rate schemes, the number of bits transmitted during a signaling interval is a random variable. A simple example of such a scheme is given in [13], where a binary data stream is parsed into the codewords of a variable-length prefix code, which are then mapped onto the points of a constellation. (A more detailed account of this type of scheme is given in [14].) Other examples of variable-rate schemes that may be interpreted as mapping the words of a prefix code onto a constellation include the shaping schemes described by Livingston [15], the block-encoded modulation schemes of Chouly and Sari [16], and the signaling schemes with "opportunistic secondary channels" described by Forney and Wei [2] (see also [17]–[21]). It should be noted, however, that these latter schemes were not constructed with shaping gain in mind.

III. Definitions

In this section, definitions for various parameters used throughout this paper are provided. Most of these parameters are carefully defined in [2] for the case of uniform signaling; here these definitions are extended to the case of nonuniform signaling.

Throughout this paper, we deal with constellations Ω embedded in an N-dimensional (ND) vector space with a well-defined (Euclidean) distance and norm. The size of Ω, $|\Omega|$,

is usually finite, but not necessarily so (as in the case of an infinite lattice). Often, Ω will be obtained as the intersection of a lattice Λ (or a translate $\mathbf{a} + \Lambda$) with a finite region $\mathbf{R}$, in which case we denote the constellation by $\Omega(\Lambda, \mathbf{R})$. The energy (squared norm) of a point $\mathbf{r} \in \Omega$ is denoted by $\|\mathbf{r}\|^2$. We shall usually assume that the transmitter produces a sequence of symbols drawn independently from the constellation and that the symbols are selected at some regular symbol rate. The probability with which the transmitter selects a point $\mathbf{r} \in \Omega$ is denoted by $p(\mathbf{r})$. As usual in the study of data transmission schemes, we are concerned with trade-offs among three parameters: reliability, bit rate, and transmitter power.

A. Reliability, Bit Rate, Transmitter Power

Although the most natural reliability measure for symbol transmission in a noisy channel is, perhaps, P_e (the average symbol error rate), this measure is often difficult to compute and to work with, especially in the case of complicated multidimensional constellations. A simple (and well-established) reliability measure for a signaling scheme on the Gaussian channel is the parameter $d^2_{\min}$, the minimum squared Euclidean distance between different constellation points. Formally,

$$d^2_{\min} \triangleq \min\{d^2(\mathbf{r}, \mathbf{r}') : \mathbf{r}, \mathbf{r}' \in \Omega, \mathbf{r} \neq \mathbf{r}'\}$$

where $d^2(\mathbf{r}, \mathbf{r}')$ denotes the squared Euclidean distance between constellation points $\mathbf{r}$ and $\mathbf{r}'$. Schemes with greater $d^2_{\min}$ will tend to have smaller symbol error rate, at least for large SNR (signal-to-noise ratio), and hence greater reliability; thus, we will take $d^2_{\min}$ as the principal reliability measure in this paper.

In fact, $d^2_{\min}$ can be used to estimate the symbol error rate at moderate to high SNR's. Let $N_{\min}(\mathbf{r})$ denote the number of constellation points at distance $d^2_{\min}$ from the point $\mathbf{r}$. The error coefficient $\overline{N}$ is the average of $N_{\min}$ over the constellation; that is,

$$\overline{N} \triangleq \sum_{\mathbf{r} \in \Omega} p(\mathbf{r}) N_{\min}(\mathbf{r}). \tag{1}$$

Assuming a white Gaussian noise channel with a one-sided noise power spectral density of N_0 W/Hz, the dominant term in a simple union bound on P_e, assuming maximum-likelihood decoding, gives us the estimate

$$P_e \approx \overline{N} Q\left(\sqrt{d^2_{\min}/2N_0}\right) \tag{2}$$

where $Q(x) \triangleq (1/\sqrt{2\pi}) \int_x^\infty \exp(-u^2/2)\, du$. It is important to note that $\overline{N}$, unlike $d^2_{\min}$, *is* affected by the probability with which constellation points are selected.

In information-theoretic terms, the transmitter is a discrete memoryless source whose output alphabet is the set of points in the constellation. The (average) bit rate is equal to the entropy of the transmitter in bits per transmitted symbol. In formal terms, a scheme with a constellation Ω, in which the symbol $\mathbf{r}$ is selected independently with probability $p(\mathbf{r})$, has bit rate

$$\beta \triangleq -\frac{2}{N} \sum_{\mathbf{r} \in \Omega} p(\mathbf{r}) \log_2 [p(\mathbf{r})] \quad \text{bits/2-D channel-use} \tag{3}$$

where, as in [2], we have normalized to 2-D. For the special case of a uniform probability distribution over a finite N-D constellation Ω, the bit rate is $(2/N) \log_2 |\Omega|$ bits/2-D channel-use, and this is the maximum bit rate for the given constellation. A signaling scheme with a normalized bit rate of β can send roughly β bits/s/Hz when implemented with a QAM (quadrature amplitude modulation) modem (which sends sequences of 2-D signals).[1]

Transmitter power is proportional to the average energy per transmitted symbol. The normalized average energy per symbol per two dimensions is given by

$$E \triangleq \frac{2}{N} \sum_{\mathbf{r} \in \Omega} p(\mathbf{r}) \|\mathbf{r}\|^2. \tag{4}$$

B. Constellation Figure of Merit and Gain

In any data transmission scheme, we would like to transmit at a large bit rate, with as high a reliability and as low a transmitter power as possible. A commonly used figure of merit for a signaling scheme, sometimes termed the "constellation figure of merit" or CFM [2], is the dimensionless, scale-invariant ratio $\text{CFM} \triangleq d^2_{\min}/E$.

To compare the relative energy efficiency of two schemes, we use the estimate (2).

Let

$$P_{e,1}(E/N_0) \approx \overline{N}_1 Q\left(\sqrt{CFM_1 E/2N_0}\right)$$

and

$$P_{e,2}(E/N_0) \approx \overline{N}_2 Q\left(\sqrt{CFM_2 E/2N_0}\right)$$

denote symbol error rate estimates for two schemes operating at the same bit rate and having, respectively, constellation figures of merit CFM_1 and CFM_2 and error coefficients $\overline{N}_1$ and $\overline{N}_2$. Then, $P^{-1}_{e,1}(p)$ and $P^{-1}_{e,2}(p)$ denote, respectively, the approximate E/N_0 value needed by each of the two schemes to achieve the symbol error rate p. At a fixed symbol error rate p, the gain, $G(p)$, of the first scheme relative to the second is given by the ratio of E/N_0 values, i.e.,

$$\begin{aligned} G(p) &= P^{-1}_{e,2}(p)/P^{-1}_{e,1}(p) \\ &= (\text{CFM}_1/\text{CFM}_2) \times \gamma_N(p, \overline{N}_1, \overline{N}_2) \end{aligned} \tag{5}$$

where

$$\gamma_N(p, \overline{N}_1, \overline{N}_2) \triangleq [Q^{-1}(p/\overline{N}_2)/Q^{-1}(p/\overline{N}_1)]^2.$$

(This latter factor is easily evaluated via a convenient approximation for $Q^{-1}(x)$ due to Hastings [22] given in Abramowitz and Stegun [23, section 26.2.23].) For small values of x,

$$Q^{-1}(x) \approx \sqrt{\ln(1/x^2)}, \tag{6}$$

and hence the asymptotic gain

$$G \triangleq \lim_{p \to 0} G(p) = \text{CFM}_1/\text{CFM}_2$$

[1] Actually, a spectral throughput of β bits per second per Hertz is achieved with QAM only by using ideal Nyquist pulse shaping. In practice, some excess bandwidth is needed, so β represents an upper bound on spectral throughput.

is the ratio of constellation figures of merit, and is not affected by error coefficient values. At nonzero values of p, the asymptotic gain will be affected by the γ_N factor in (5). In particular, when $\overline{N}_1 > \overline{N}_2$, the asymptotic gain will be reduced. When p is small, we can again use (6) and write

$$\gamma_N(p, \overline{N}_1, \overline{N}_2) \approx [1 + \log(\overline{N}_1/\overline{N}_2)/\log(\overline{N}_2/p)]^{-1} \quad (7)$$

to estimate this effect.

For bit rates $\beta \geq 2$ bits/2-D channel-use, we take as a baseline for comparison the CFM obtainable with a simple 1-D PAM (pulse amplitude modulation) constellation. Assuming that constellation points are selected with equal probability, this baseline has figure of merit given by $\mathrm{CFM}_\oplus = 6/(2^\beta - 1)$. The (asymptotic) gain of a scheme with a given CFM, relative to the baseline, is then given by

$$G(d^2_{\min}, \beta, E) \triangleq \mathrm{CFM}/\mathrm{CFM}_\oplus = \frac{(2^\beta - 1)\, d^2_{\min}}{6E}, \qquad \beta \geq 2. \quad (8)$$

This gain measure effectively combines our reliability, bit rate, and transmitter power measures into a single figure, thus allowing for direct comparisons of disparate schemes. Note also that the gain formula (8) is valid only for $\beta \geq 2$, i.e., for bit rates of at least one bit per symbol per dimension. Since we wish to study bandwidth-efficient schemes, i.e., those with large β, this restriction will pose no problem, and we assume it to hold throughout this paper.

While we will be interested primarily in asymptotic gain in this paper, we can use (5) to estimate the effect of the error coefficient at nonzero error probability values. In Appendix A, we show that a cubic N-D constellation based on $\mathbf{Z}^N$ under uniform signaling and supporting a bit rate β (i.e., an N-D baseline constellation) has error coefficient $\overline{N}_\oplus = 2N(1 - 2^{-\beta/2})$. Thus, in two dimensions and for large values of β, $\overline{N}_\oplus \approx 4$. Using (7), we estimate that every factor of two increase in the error coefficient (over the baseline of 4) reduces the asymptotic gain by about 0.2 dB for p on the order of 10^{-6}, in agreement with the rule of thumb given by Forney [10, p. 1142]. (A plot of γ_N in dB versus $\log_2 \overline{N}$ suggests that a loss of 0.22 dB per error coefficient doubling is a fairly accurate rule, up to about $\overline{N} = 64$, at $p = 10^{-6}$. Similarly, at $p = 10^{-8}$, the loss is about 0.17 dB per doubling in error coefficient.)

It is important to note that when the inner points of a constellation are selected more often than the outer points, the error coefficient will increase, because the inner points tend to have greater $N_{\min}$ than the outer points. For example, consider a constellation drawn from the $\mathbf{Z}^2$ lattice, for which $\overline{N} \leq 4$. In light of the previous paragraph, we can estimate the maximum loss in gain by

$$0.22 \log_2 4/N_\oplus \; dB = -0.22 \log_2 (1 - 2^{-\beta/2})\; dB$$

for error rates on the order of 10^{-6}. For $\beta \approx 2$, when the baseline has $\overline{N} = 2$, this maximum possible degradation is quite large (0.22 dB) relative to the maximum achievable shaping gain (1.53 dB); however, for $\beta \approx 6$, the maximum degradation is only about 0.04 dB. At smaller error rates, the degradation is even less. As mentioned, we prefer to focus on the asymptotic gain, which is unaffected by variations in the error coefficient, but we caution the reader to note that the error coefficient must be accounted for in estimating the gain at nonzero error probabilities.

Suppose now that the constellation Ω is obtained from an infinite ND lattice Λ, and that a point $\boldsymbol{r} \in \Omega$ is selected with probability $p(\boldsymbol{r})$. If $\Omega \neq \Lambda$, it is convenient to extend the distribution $p(\boldsymbol{r})$ to all the points of Λ, simply by assigning zero probability to any points not in Ω. Let $V(\Lambda)$ be the volume of a fundamental region [24] of Λ, i.e., the volume of N-space associated with each lattice point, and let $2^{N\beta(p)/2}V(\Lambda)$ be the "entropy volume" of Λ with distribution $p(\boldsymbol{r})$. We can then write $G(d^2_{\min}, \beta, E)$ (8) as

$$G(d^2_{\min}, \beta, E) = \gamma_c(\Lambda)\gamma_s(p)(1 - 2^{-\beta(p)}), \quad (9)$$

where

$$\gamma_c(\Lambda) = d^2_{\min}(\Lambda)/V(\Lambda)^{2/N} \quad (10)$$

is the coding gain [10] of the lattice Λ, and

$$\gamma_s(p) = 2^{\beta(p)}V(\Lambda)^{2/N}/[6E(p)]$$

is the shaping gain of Λ with distribution $p(\boldsymbol{r})$. For large bit rates β, when the discretization factor $\gamma_d(p) = 1 - 2^{-\beta(p)}$ is small, the total gain is approximately separable into the product of a coding gain and a shaping gain. The coding gain $\gamma_c(\Lambda)$, a geometric property of the lattice Λ studied in the coded modulation literature and elsewhere [10], [11], [24], is independent of the probability distribution $p(\boldsymbol{r})$ used to select constellation points and therefore not of central interest in this paper. The shaping gain $\gamma_s(p)$ is largely independent of the underlying lattice Λ, except insofar as the lattice restricts the distribution $p(\boldsymbol{r})$; this is why we have suppressed an explicit dependence on Λ in our notation. In general [2], $\gamma_s(p) \lesssim \pi e/6$. As we shall see, by choosing the distribution $p(\boldsymbol{r})$ to be the Maxwell–Boltzmann distribution, $\gamma_s(p)$ can be made to approach the ultimate shaping gain of $\pi e/6$ in *any* dimension.

C. Other Constellation Parameters

Often, higher-dimensional constellations Ω are obtained as subsets drawn from Cartesian products of lower-dimensional "constituent" constellations. When the dimension N of Ω is even, we may define the constituent 2-D constellation of Ω as the smallest 2-D constellation Ω_2 such that $\Omega \subset \Omega_2^{N/2}$, where $\Omega_2^{N/2}$ is the $N/2$-fold Cartesian product of Ω_2 with itself. The constituent 2-D constellation Ω_2 plays an important role when the signaling scheme is to be implemented with a QAM modem, since all transmitted signals are obtained as sequences of QAM signals drawn from Ω_2.

If Ω is odd-dimensional, then Ω^2 is even-dimensional and can be implemented with a QAM modem. This suggests that we may define the constituent 2-D constellation of Ω as the constituent 2-D constellation of Ω^2. In particular, if $N = 1$, this implies $\Omega_2 = \Omega^2$. For later use, we note that the constituent 2-D constellation of $\mathbf{B}_N(R)$, an N-ball of radius R centered at the origin, is a 2-D disk $\mathbf{B}_2(R)$ when N is even, and is a square $\mathbf{B}_1^2(R)$ of side $2R$ when N is odd.

As discussed in [2] and [25], an important parameter in the design of a signaling scheme is the 2-D "constellation expansion ratio"

$$\mathrm{CER2}(\Omega) \triangleq |\Omega_2|/2^{\beta},$$

where $|\Omega_2|$ is the number of points in the constituent 2-D constellation of Ω, and 2^{β} represents the number of points in a comparable baseline constellation supporting the same bit rate with uniform signaling. Since $|\Omega_2| \geq |\Omega|^{2/N} \geq 2^{\beta}$, we have $\mathrm{CER2}(\Omega) \geq 1$. Large 2-D constellations are sensitive to nonlinearities and other signal-dependent perturbations, so it is desirable that $\mathrm{CER2}(\Omega)$ be as close to its lower bound of unity as possible.

Another important parameter discussed in [2] is the 2-D "peak-to-average energy ratio"

$$\mathrm{PAR2}(\Omega) \triangleq r_{\max}^2(\Omega_2)/E, \tag{11}$$

where $r_{\max}^2$ is the energy maximum of the points in Ω_2, the constituent 2-D constellation of Ω, and E is the normalized average energy. PAR2 is a measure of the dynamic range of the signals transmitted by a QAM modem. To minimize the effects of signal-dependent distortion, it is desirable that PAR2, like CER2, be as small as possible. (Note that the "peak" energy in this definition is found by averaging a signal over a 2-D interval, thus making PAR2 independent of the pulse shape used in implementation. The actual "instantaneous" peak energy depends on the actual pulses used, and on how these pulses superpose in time when transmitted in sequence.)

Since constellations are often obtained by taking the intersection of an N-D lattice Λ with an N-D region $\mathbf{R}$, in analogy with constituent 2-D constellations, it is useful to define constituent 2-D lattices and regions. We define the constituent 2-D lattice Λ_2 of Λ as the smallest 2-D set of Λ_2 such that $\Lambda^2 \subset \Lambda_2^N$. Similarly the constituent 2-D region of $\mathbf{R}$ is the smallest 2-D region $\mathbf{R}_2$ such that $\mathbf{R}^2 \subset \mathbf{R}_2^N$. If $N = 1$, then $\Lambda_2 = \Lambda^2$ and $\mathbf{R}_2 = \mathbf{R}^2$. It is clear that Ω_2 is obtained as the intersection of Λ_2 with $\mathbf{R}_2$.

IV. The Maxwell–Boltzmann Distribution

As pointed out in the Introduction, it follows immediately from the maximum entropy principle that the Maxwell–Boltzmann distribution maximizes bit rate for a fixed average energy. (For a good introduction to the maximum entropy principle, see [1, ch. 11].) Equivalently, the Maxwell–Boltzmann distribution minimizes average energy for a fixed bit rate.[2] Signal point selection with a Maxwell–Boltzmann distribution causes a constellation point $\boldsymbol{r}$, with energy $\|\boldsymbol{r}\|^2$, to be selected with probability $p(\boldsymbol{r}) \propto \exp(-\lambda\|\boldsymbol{r}\|^2)$, where the parameter $\lambda \geq 0$ governs the trade-off between bit rate and average energy. More precisely, the optimal distribution is one in which

$$p(\boldsymbol{r}) \triangleq \exp(-\lambda\|\boldsymbol{r}\|^2)/Z(\lambda), \qquad \lambda \geq 0,$$

where the partition function $Z(\lambda)$ is chosen to normalize the distribution, i.e.,

$$Z(\lambda) \triangleq \sum_{\boldsymbol{r}\in\Omega} \exp(-\lambda\|\boldsymbol{r}\|^2), \qquad \lambda \geq 0. \tag{12}$$

(For infinite constellations, λ must be strictly positive.) The Maxwell–Boltzmann distribution arises in many contexts; e.g., in the optimization of permutation modulation for quantization [26] and for transmission [27], in neural networks [28], and in simulated annealing [29], among others.

For finite constellations, setting $\lambda = 0$ yields a signaling scheme in which constellation points are selected with uniform probability; thus, "classical" fixed-rate signaling schemes appear here as a special case. Note too that, with a Maxwell–Boltzmann distribution, outer points (points with large energy) are never selected more often than inner points (points with small energy). An equivalence class of the points of Ω all having the same energy is called a *shell* of the constellation. With a Maxwell–Boltzmann distribution, the points of a shell are selected equally often. Indeed, if all constellation points lie in the same shell, "classical" uniform signaling is obtained for all values of λ.

In statistical mechanics, much attention is paid to the computation of the partition function $Z(\lambda)$ (12) in various physical systems. This is due to the fact that the average energy and entropy are easily obtained in terms of $Z(\lambda)$. Indeed, the normalized average energy (4) is obtained as

$$E(\lambda) = \frac{2}{N}\left(\frac{-d\ \ln\ Z(\lambda)}{d\lambda}\right), \tag{13}$$

and the normalized bit rate (3) is obtained as

$$\beta(\lambda) = \frac{2}{N}\left[-\lambda^2\frac{d}{d\lambda}\left(\frac{\log_2 Z(\lambda)}{\lambda}\right)\right] = \frac{2}{N}\log_2 Z(\lambda) + \frac{\lambda E(\lambda)}{\ln 2}. \tag{14}$$

The partition function is easily obtained in terms of the theta series [30] or Euclidean weight distribution [10] of a constellation. The theta series for a constellation Ω is simply a generating function for the set of energy values (squared norms) taken on by the points of Ω, and is defined as $\Theta(x) \triangleq \sum_{\boldsymbol{r}\in\Omega} x^{\|\boldsymbol{r}\|^2}$, where we interpret $\Theta(x)$ as a real function of x, and note that $Z(\lambda) = \Theta[\exp(-\lambda)]$.

The Maxwell–Boltzmann parameter λ governs the trade-off between bit rate and average energy. In analogy with statistical mechanics, we might call λ the "inverse temperature" of the Maxwell–Boltzmann distribution, i.e., $\lambda = 1/(kT)$, where, in statistical mechanics, k is the Boltzmann constant and T is the temperature. When $\lambda = 0$ (infinite "temperature"), the uniform distribution is obtained, corresponding to the maximum possible entropy for the given constellation. (In statistical mechanics, all states of a system are equally occupied at infinite temperature.) As $\lambda \to \infty$ (or the "temperature" cools toward absolute zero), the bit rate as well as the average energy are reduced as the points with large energy are selected less frequently. The "limiting constellation" (obtained at absolute zero "temperature") consists of only the innermost points

[2] In both optimization problems, the given constellation must be able to *support* the given bit rate or the given average energy, so these values are themselves constrained; this is pointed out at the end of Section IV.

of the original constellation (the ground states in statistical mechanics), and these points are selected equally often.

An important property of the Maxwell–Boltzmann distribution is its "separability" property. Suppose the ND constellation Ω is the Cartesian product of two or more "factor constellations," i.e., $\Omega = \Omega_1 \times \Omega_2 \times \cdots \times \Omega_J$, $J \geq 2$, where Ω_i is N_i-dimensional and $\sum_{i=1}^{J} N_i = N$. Then it follows that

$$Z(\lambda) = \prod_{i=1}^{J} Z_i(\lambda), \tag{15}$$

where $Z_i(\lambda)$ is the partition function over the ith factor constellation, i.e., $Z_i(\lambda) = \sum_{r_i \in \Omega_i} \exp(-\lambda \|r_i\|^2)$. When (15) holds, the Maxwell–Boltzmann distribution with parameter λ is separable into the product of Maxwell–Boltzmann distributions over the factor constellations, each with parameter λ. In practice, this means that optimal nonuniform signaling can be implemented on separable constellations by independently implementing nonuniform signaling on each of the factor constellations. From (13) and (14) it follows that E and β can be obtained by a weighted average of the corresponding factor constellation quantities, i.e., $E = \sum_{i=1}^{J} E_i N_i / N$ and $\beta = \sum_{i=1}^{J} \beta_i N_i / N$.

If the constellation Ω has $|\Omega|$ points in total, and N_{in} "innermost" points of minimum energy, then as λ ranges from 0 to $+\infty$, every value of β in the range $(2 \log_2 (N_{in})/N, 2 \log_2 (|\Omega|)/N]$ is obtained, and we say that the constellation supports every bit rate in this range with a Maxwell–Boltzmann distribution. Similarly, over the same range of λ values, the constellation supports average energy values in the range $(E_{\min}, E_{\text{uniform}}]$, where $E_{\min}$ is the normalized energy of a minimum-energy point in Ω, and E_{uniform} is the normalized average energy under uniform signaling. An infinite lattice can support any positive bit rate and any positive average energy with a Maxwell–Boltzmann distribution.

V. Spherical Constellations

In this section, we apply the Maxwell–Boltzmann distribution to spherical constellations based on the densest known lattices in various numbers N of dimensions. In one dimension, the constellations are based on the integer lattice $\mathbf{Z}$, while in two dimensions, the constellations are based on A_2, the hexagonal lattice. In higher dimensions, the constellations are based on the lattices denoted D_4, E_6, E_8, and K_{12}; the subscript in this notation displays the number of dimensions N. Basis vectors and extensive theta series tables for these lattices are given in [24, ch. 4]. Each constellation is obtained by taking some number M of the points of smallest energy in the lattice, where M is chosen so as to include some integral number of lattice shells. Equivalently, we may think of the constellations as being obtained by forming the intersection of the infinite lattice with an N-sphere centered at the origin; hence the term spherical constellations.[3] We note that, due to the separability property of the Maxwell–Boltzmann distribution, the results we obtain for these spherical constellations are also applicable to those nonspherical constellations that can be expressed as Cartesian products of spherical constellations. For example, N-cube shaped constellations based on the integer lattice $\mathbf{Z}^N$ are Cartesian products of simple 1-D "spherical" constellations based on $\mathbf{Z}$.

[3] Caution: some authors reserve this term to refer to constellations in which all points are on the surface of a sphere; in this paper, spherical constellations do, in general, include interior points as well.

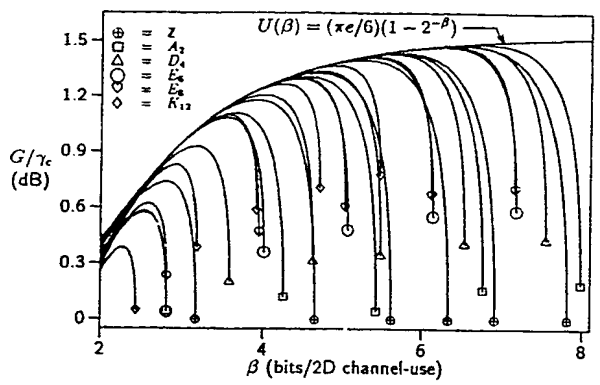

Fig. 2. Normalized gain of spherical constellations drawn from various dense lattices with signal point selection performed according to the Maxwell–Boltzmann distribution.

Plotted in Fig. 2 is the "normalized" gain that selected constellations provide for various bit rates when constellation points are selected with a Maxwell–Boltzmann distribution. The normalized gain is obtained from the gain G, defined in (8), by dividing by the coding gain $\lambda_c(\Lambda)$ (10), of the lattice from which the constellation is drawn. Each curve in Fig. 2 is obtained by varying the parameter λ from zero—corresponding to the maximum bit rate (rightmost) point in each curve—through positive values. (Recall that $\lambda = 0$ corresponds to the "classical" case of a uniform distribution.) Curves corresponding to constellations drawn from the same lattice but extending further to the right, i.e., to larger bit rates, correspond to larger constellations. Also plotted in Fig. 2 is the function

$$U(\beta) \triangleq (\pi e/6)(1 - 2^{-\beta}). \tag{16}$$

As we shall explain in Section VI, for bit rates β greater than about 2.5, $U(\beta)$ forms the "upper envelope" of the gain curves, to good approximation.

Fig. 2 has several noteworthy features. First, notice that the gain obtained by nonuniform signaling with a constellation can significantly exceed that provided under uniform signaling, at the expense of a reduction in bit rate. The additional gain provided under nonuniform signaling is called the "biasing gain" [7].

The curves corresponding to large constellations tend to merge with curves corresponding to smaller constellations as λ is increased. This happens because the smaller constellations are subconstellations of the large constellations. As λ is increased, the outer points of the large constellations are selected very infrequently, so that, in effect, these outer points can be neglected and the large constellation "shrinks" into a smaller constellation.

Note also that each curve tends to merge with the $U(\beta)$ curve. Comparing (16) to (9), we see that this merging implies that the shaping gain under nonuniform signaling approaches the ultimate limit of $\pi e/6$ as λ becomes large, and that this limit is obtained independently of dimension.

Fig. 2 also illustrates the "law of diminishing returns" governing the biasing gain. Recall that the rightmost point of each curve ($\lambda = 0$) in the graph corresponds to uniform signaling. We see that, for constellations having dimension greater than unity, some "initial" gain is available under uniform signaling. Furthermore, the initial gain increases with increasing dimension. This initial gain is, of course, the shaping gain of the N-sphere under uniform signaling, which increases with N and ultimately approaches the value $\pi e/6$. This forces the ultimate biasing gain (the difference between the ultimate shaping gain and the shaping gain of the N-sphere) to decrease with dimension.

Qualitatively, this law of diminishing returns arises due to a phenomenon known as the "sphere hardening effect" (see, e.g., [31]). In a many-dimensional sphere, almost all of the volume is located near the surface of the sphere; consequently, almost all constellation points lie near or on the surface as well. Since these points all have the same energy, i.e., the same cost, signaling with a Maxwell–Boltzmann distribution will cause these points to be selected equally often. Thus, uniform signaling with spherical constellations becomes increasingly effective as the dimension increases, ultimately approaching the performance of nonuniform signaling.

VI. Continuous Approximations

Let $\Omega(\Lambda, \mathbf{R})$ be a constellation obtained from the intersection of an N-D lattice Λ (or a translate $\boldsymbol{a} + \Lambda$ of Λ) with a finite N-D region $\mathbf{R}$. In [2], Forney and Wei were able to obtain much insight into the performance of such constellations via the so-called "continuous approximation," obtained by replacing discrete sums over the points of Ω with properly normalized integrals over the region $\mathbf{R}$. In this section we use the same approach to obtain similar insight in the case of nonuniform signaling, essentially by replacing discrete Maxwell–Boltzmann distributions with continuous Gaussian distributions, truncated to the region $\mathbf{R}$. The continuous approximation allows us to obtain a continuous approximation for the partition function $Z(\lambda)$, from which estimates of all other relevant system parameters are obtained. The result of Forney and Wei for uniform signaling [2] appear as a special case, obtained when the Maxwell–Boltzmann parameter λ is set to zero.

A. Energy and Entropy Approximations

Let f: $\mathbf{R}^N \rightarrow \mathbf{R}$, a function of N variables, be Riemann-integrable over an ND region $\mathbf{R}$. Given an ND lattice Λ^*, we have

$$\int_{\mathbf{R}} f(\boldsymbol{r})\, dV(\boldsymbol{r}) = \lim_{\alpha \rightarrow 0} \sum_{\boldsymbol{r} \in \Omega(\alpha\Lambda^*, \mathbf{R})} f(\boldsymbol{r}) V(\alpha\Lambda^*), \tag{17}$$

where $\Omega(\alpha\Lambda^*, \mathbf{R}) \triangleq \alpha\Lambda^* \cap \mathbf{R}$, and $V(\alpha\Lambda^*)$ denotes the volume of a fundamental region of the scaled lattice $\alpha\Lambda^*$. For small α, we expect the summation on the right-hand side of (17) to be a good approximation to the integral on the left-hand side of (17). Setting $\Lambda = \alpha\Lambda^*$ (α small), we obtain from (17) the general continuous approximation

$$\sum_{\boldsymbol{r} \in \Omega(\Lambda, \mathbf{R})} f(\boldsymbol{r}) \approx V(\Lambda)^{-1} \int_{\mathbf{R}} f(\boldsymbol{r})\, dV(\boldsymbol{r}). \tag{18}$$

For example, setting $f(\boldsymbol{r}) = 1$, yields the approximation

$$|\Omega(\Lambda, \mathbf{R})| = \sum_{\boldsymbol{r} \in \Omega(\Lambda, \mathbf{R})} 1 \approx V(\mathbf{R})/V(\Lambda), \tag{19}$$

where $V(\mathbf{R})$ denotes the volume of the region $\mathbf{R}$. This is "Proposition 1" of Forney and Wei [2].

For nonuniform signaling with a Maxwell–Boltzmann distribution, the average energy and bit rate are determined by the partition function $Z(\lambda)$ (12). The continuous approximation (18) yields

$$Z(\lambda) \approx V(\Lambda)^{-1} \int_{\mathbf{R}} \exp(-\lambda \|\boldsymbol{r}\|^2)\, dV(\boldsymbol{r}). \tag{20}$$

Combining approximation (20) with (13) we approximate E, the normalized average energy (4), by

$$E(\mathbf{R}, \lambda) \triangleq \frac{2}{N} \int_{\mathbf{R}} \|\boldsymbol{r}\|^2 f(\boldsymbol{r}, \lambda)\, dV(\boldsymbol{r}), \tag{21}$$

where

$$f(\boldsymbol{r}, \lambda) = \frac{\exp(-\lambda \|\boldsymbol{r}\|^2)}{\int_{\mathbf{R}} \exp(-\lambda \|\boldsymbol{r}\|^2)\, dV(\boldsymbol{r})}. \tag{22}$$

Note that $f(\boldsymbol{r}, \lambda)$ represents a continuous Gaussian probability density function, truncated to the region $\mathbf{R}$. The continuous approximation (21) estimates the normalized average constellation energy by the normalized average energy of this continuous random variable.

In the same way, combining approximation (20) with (14), we approximate β, the normalized bit rate (3), by

$$\beta(\mathbf{R}, \lambda) \triangleq \frac{2}{N}[H(\mathbf{R}, \lambda) - \log_2 V(\Lambda)], \tag{23}$$

where

$$H(\mathbf{R}, \lambda) = -\int_{\mathbf{R}} f(\boldsymbol{r}, \lambda) \log_2 f(\boldsymbol{r}, \lambda)\, dV(\boldsymbol{r})$$

is the differential entropy of a continuous Gaussian random variable, truncated to the region $\mathbf{R}$. The continuous approximations (21) and (23) were also used by Forney and Wei [2, section IV-B], and by Forney in [6, section V].

B. Shaping Gain Approximation

We know use (21) and (23) to estimate the gain provided by nonuniform signaling. Writing $\beta(\mathbf{R}, \lambda)$ for β, and $E(\mathbf{R}, \lambda)$ for E, in the gain expression (8) yields the approximation $G \approx G(\mathbf{R}, \lambda)$ where

$$\begin{aligned} G(\mathbf{R}, \lambda) &\triangleq \frac{2^{\beta(\mathbf{R}, \lambda)} d_{\min}^2}{6E(\mathbf{R}, \lambda)} (1 - 2^{-\beta(\mathbf{R}, \lambda)}) \\ &= \left[\frac{d_{\min}^2}{V(\Lambda)^{2/N}}\right] \cdot \left[\frac{2^{(2/N)H(\mathbf{R}, \lambda)}}{6E(\mathbf{R}, \lambda)}\right] \\ &\quad \times (1 - 2^{-\beta(\mathbf{R}, \lambda)}). \end{aligned} \tag{24}$$

As in (9), we have grouped the gain $G(\mathbf{R}, \lambda)$ (24) into three factors—the first being the coding gain $\gamma_c(\Lambda) \triangleq d^2_{\min}/V(\Lambda)^{2/N}$ of the lattice Λ [10], the second being a continuous approximation for the shaping gain, namely,

$$\gamma_s(\mathbf{R}, \lambda) \triangleq \frac{2^{(2/N)H(\mathbf{R}, \lambda)}}{6E(\mathbf{R}, \lambda)}, \tag{25}$$

and the third being the discretization factor

$$\gamma_d(\mathbf{R}, \lambda) \triangleq 1 - 2^{-\beta(\mathbf{R}, \lambda)}.$$

The discretization factor was omitted in the analysis of Forney and Wei [2], who were interested in asymptotic ($\beta \to \infty$) limits for the shaping gain. However, for most practical values of β, this factor is not insignificant and should not be omitted. Indeed, as will become evident, including this factor provides accurate estimates for gain, even for relatively small values of β (see Fig. 3 and Section VI-D).

We now estimate the biasing gain [7], i.e., the additional gain that nonuniform signaling can provide over uniform signaling with the same constellation. Setting $\lambda = 0$ reduces (25) to the special case of uniform signaling discussed in [2]. The shaping gain $\gamma_s(\mathbf{R}, \lambda)$ can be written in terms of $\gamma_s(\mathbf{R}, 0)$ as

$$\begin{aligned}\gamma_s(\mathbf{R}, \lambda) &= \gamma_s(\mathbf{R}, 0) \cdot \left[\frac{2^{(2/N)H(\mathbf{R}, \lambda)}}{2^{(2/N)H(\mathbf{R}, 0)}}\right] \cdot \left[\frac{E(\mathbf{R}, 0)}{E(\mathbf{R}, \lambda)}\right] \\ &= \gamma_s(\mathbf{R}, 0) \cdot \left[\frac{2^{\beta(\mathbf{R}, \lambda)}}{2^{\beta(\mathbf{R}, 0)}}\right] \cdot \left[\frac{E(\mathbf{R}, 0)}{E(\mathbf{R}, \lambda)}\right] \\ &= \gamma_s(\mathbf{R}, 0) \cdot 2^{-\rho(\mathbf{R}, \lambda)} \cdot g_E(\mathbf{R}, \lambda), \end{aligned} \tag{26}$$

where

$$\rho(\mathbf{R}, \lambda) \triangleq \beta(\mathbf{R}, 0) - \beta(\mathbf{R}, \lambda) \tag{27}$$

is the normalized redundancy (or loss in bit rate) caused by selecting constellation points with a nonuniform distribution. Clearly, $\gamma_b(\mathbf{R}, \lambda)$, the total biasing gain, is given by the product of the second and third factors in (26), i.e.,

$$\gamma_b(\mathbf{R}, \lambda) = g_E(\mathbf{R}, \lambda) 2^{-\rho(\mathbf{R}, \lambda)}. \tag{28}$$

In (28), we have identified two separate factors that characterize the biasing gain. The "energy savings factor"

$$g_E(\mathbf{R}, \lambda) \triangleq E(\mathbf{R}, 0)/E(\mathbf{R}, \lambda) \geq 1, \tag{29}$$

accounts for the energy savings that result when constellation points of low energy are selected more often that points of large energy. Of course, selecting points with a nonuniform distribution results in a loss of entropy and hence a drop in the baseline average energy. The "energy loss factor" $2^{-\rho(\mathbf{R}, \lambda)}$ accounts for this drop.

C. CER2 and PAR2 Approximations

We now provide continuous approximations for CER2 and PAR2, two constellation parameters defined in Section III. From (19), we estimate $|\Omega_2| \approx V(\mathbf{R}_2)/V(\Lambda_2)$ for the size of the constituent 2-D constellation, where $V(\mathbf{R}_2)$ is the volume (actually area) of $\mathbf{R}_2$, the constituent 2-D region of $\mathbf{R}$, and $V(\Lambda_2)$ is the fundamental volume (actually area) o Λ_2, the constituent 2-D lattice of Λ. Combining this with ou approximation (23) for the bit rate, we obtain

$$\mathrm{CER2}(\Omega) \approx [V(\Lambda)^{2/N}/V(\Lambda_2)] \cdot [V(\mathbf{R}_2)/2^{(2/N)H(\mathbf{R}, \lambda)}].$$

Here, as in [2], we identify two independent factors: the coding constellation expansion ratio of Λ, $\mathrm{CER2}_c(\Lambda) \triangleq V(\Lambda)^{2/N}/V(\Lambda_2)$ and the shaping constellation expansion ratio, $\mathrm{CER2}_s(\mathbf{R}, \lambda) \triangleq V(\mathbf{R}_2)/2^{(2/N)H(\mathbf{R}, \lambda)}$ Again, as in the case of gain, one component, the coding constellation expansion factor, is a geometric property of the lattice Λ and is unaffected by the probability distribution with which the constellation points are selected. The other component, the shaping constellation expansion ratio, depends both on the region $\mathbf{R}$ and the parameter λ.

When $\lambda = 0$, we obtain the special case of uniform signaling, where

$$\mathrm{CER2}_s(\mathbf{R}, 0) = V(\mathbf{R}_2)/V(\mathbf{R})^{2/N}.$$

We may write $\mathrm{CER2}_s(\mathbf{R}, \lambda)$ in terms of $\mathrm{CER2}_s(\mathbf{R}, 0)$ as

$$\mathrm{CER2}_s(\mathbf{R}, \lambda) = \mathrm{CER2}_s(\mathbf{R}, 0) \cdot 2^{\rho(\mathbf{R}, \lambda)},$$

where $\rho(\mathbf{R}, \lambda)$ (27) is the normalized redundancy under nonuniform signaling. We see that in addition to the constellation expansion due to uniform signaling, we have incurred an additional constellation expansion factor due to the loss in rate caused by nonuniform signaling. Thus, for a fixed constellation, $\mathrm{CER2}_s(\mathbf{R}, 0)$, the shaping constellation expansion ratio induced by uniform signaling, is a lower bound to $\mathrm{CER2}_s(\mathbf{R}, \lambda)$, the shaping constellation expansion ratio under nonuniform signaling.

To estimate $\mathrm{PAR2}(\Omega)$ (11), we note that the peak energy of the constituent 2-D constellation is a geometric property unaffected by the probability with which constellation points are selected. The average energy, on the other hand, is reduced from its value under uniform signaling by the energy savings factor $g_E(\mathbf{R}, \lambda)$; thus PAR2 is increased by the same factor, i.e.,

$$\mathrm{PAR2}(\mathbf{R}, \lambda) = \mathrm{PAR2}(\mathbf{R}, 0) g_E(\mathbf{R}, \lambda),$$

where $\mathrm{PAR2}(\mathbf{R}, 0)$ denotes the PAR2 under uniform signaling.

D. Applying the Continuous Approximations

In applying these continuous approximations to actual constellations, one is confronted with a certain flexibility in the choice of approximating region $\mathbf{R}$. For example, to approximate the behavior of an M point, symmetric PAM constellation based on $\mathbf{Z}$, one would choose $\mathbf{R} = [-R, R]$, a 1D "sphere" of radius R; however, it is not clear which choice for radius R is best. Indeed, the "best" choice for R depends both on the constellation parameter—be it average energy, entropy, or whatever—that one is trying to estimate, and on the value of the Maxwell–Boltzmann parameter λ. The same flexibility in choice of sphere radius R occurs when one attempts to approximate the behavior of an N-D spherical constellation with an N-sphere $\mathbf{B}_N(R)$.

Since the size, $|\Omega|$, of the constellation is assumed known, one approach is to choose the radius R so as to match the bit rate estimate (23) at $\lambda = 0$ with the actual bit rate $(2/N)\log_2|\Omega|$ at $\lambda = 0$. In effect, this forces the volume of the region $\mathbf{R}$ to satisfy $V(\mathbf{R}) = |\Omega|V(\Lambda)$, so that (19) is satisfied with equality. From extensive numerical calculations, we have found that this approach gives very satisfactory estimates for bit rate and average energy, over a fairly wide range for λ.

Applying this approach to the M-PAM example, we find that $R = M/2$ and $E \approx M^2/6$. However, the actual energy $E = (M^2 - 1)/6$; thus, using the continuous approximation, we have incurred an error that is a factor of $(1 - M^{-2}) = 1 - 2^{-\beta}$. The discretization factor, $\gamma_d(\beta) = 1 - 2^{-\beta}$, can thus be interpreted as a correction factor used to adjust the average energy when applying the continuous approximation to the baseline constellations under uniform signaling.

E. Spherical and Cubic Constellations

Continuous approximations for the various constellation parameters are derived in Appendix B for the important case in which the region $\mathbf{R} = \mathbf{B}_N(R)$, an N-ball of radius R centered at the origin. These include cubic constellations, which are Cartesian products of 1-D "spheres," as a special case.

In order to compare the shaping gain predicted by these continuous approximations to actual shaping gain, we have plotted in Fig. 3 normalized gain curves for spherical constellations Ω in various dimensions. As in Fig. 2, the gain values are normalized by dividing the total gain by the coding gain of the lattice from which the constellation is drawn. Solid curves represent the actual normalized gain G/γ_c, computed from (8). The dotted curves give $(|\Omega|^{2/N}2^{-\rho(\mathbf{B}_N(R),\lambda)} - 1)/[6E(\mathbf{B}_n(R), \lambda)]$, as defined in Appendix B. The radius R in each case is chosen so that $V(\mathbf{B}_N(R))/V(\Lambda) = |\Omega|$. Also plotted in Fig. 3 is the function $U(\beta)$ (16) which represents the "upper envelope" shown in the figure.

We see that the curves corresponding to the approximate normalized gain closely match the actual normalized gain curves for all values of $\beta \geq 2$, although some difference is seen for small β. For large β, however, the curves corresponding to the approximation correspond with the actual normalized gain curves, confirming the asymptotic accuracy of the continuous approximation.

As pointed out in Appendix B, for large bit rates, the shaping gain under nonuniform signaling approaches $\pi e/6$, independently of dimension. The ultimate biasing gain approaches $\pi e/[6\gamma_\otimes(N)]$, where $\gamma_\otimes(N)$ (39) denotes the shaping gain of the N-sphere under uniform signaling. Since $\gamma_\otimes(N) \to \pi e/6$ monotonically from below as $N \to \infty$, γ_b approaches unity as the dimension increases, thus confirming the "law of diminishing" returns discussed at the end of Section V.

Curves showing the trade-offs between $\mathrm{CER2}_s$ and shaping gain or between PAR2 and shaping gain for ND spherical constellations are easily obtained from the continuous approximations derived in Appendix B. However, as asserted by Forney and Wei [2], the best possible trade-offs are achieved by 2-D spherical constellations, i.e., by regions shaped as discs in two dimensions. Recall that Cartesian products of

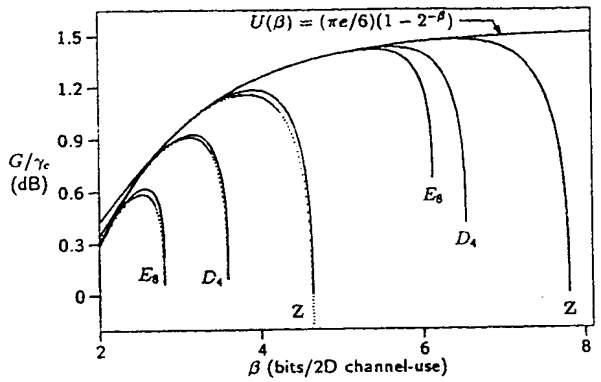

Fig. 3. Comparison of the approximate and actual normalized gain of spherical constellations based on various N-dimensional lattices. The dotted curves are obtained by applying continuous approximations for gain.

basic regions achieve the same performance as the basic region itself. Thus, the best region $\mathbf{R}$ for use with nonuniform signaling in $2n$ dimensions is the n-fold Cartesian product of a 2-D disc—a so-called polydisc [32]—because this region will achieve a given value of shaping gain with least $\mathrm{CER2}_s$ and PAR2. Thus, while nonuniform signaling will always cause a constellation expansion relative to uniform signaling with the same constellation, this constellation expansion is never greater and usually *less* than would be required under uniform signaling to achieve the same shaping gain.

VII. Shaping with Binary Prefix Codes

In this section, we study methods of achieving nonuniform signaling schemes for the transmission of binary data. Assuming, as usual, that we wish to transmit the output of a memoryless binary equiprobable source, the most obvious means of generating events with nonuniform probabilities is to parse the output of the source into codewords of variable length. Since the probability of occurrence of a codeword of length l_i is 2^{-l_i}, shorter (more frequently occurring) codewords may be mapped to constellation points with low energy and longer (less frequently occurring) codewords may be mapped to points with high energy and, in this way, shaping gain may be achieved. This approach was suggested by a brief example in [13] and is discussed in greater detail in [14].

To ensure unique and complete parsing, it can be shown (e.g., [33, p. 297]) that the variable length binary codewords must form a complete binary prefix code, in which the M codeword lengths l_i, $i = 1, \cdots, M$ satisfy

$$\sum_{i=1}^{M} 2^{-l_i} = 1. \tag{30}$$

Although we will not always explicitly refer to them as such, all prefix codes considered in this paper are complete.

A. Matched Codes

The simplest example of the idea of mapping a prefix code to a constellation is probably the following. The output of a binary equiprobable memoryless source (the data source) is parsed into a sequence of blocks drawn from the set

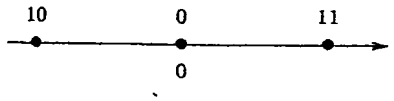

Fig. 4. A simple nonuniform PAM scheme.

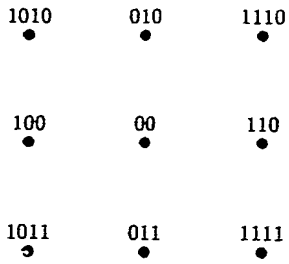

Fig. 5. A nonuniform QAM scheme. $\beta = 3$.

$\{0, 10, 11\}$ and these output blocks are mapped onto a PAM constellation as shown in Fig. 4. Since $P[0] = \frac{1}{2}$ and $P[10] = P[11] = \frac{1}{4}$, this scheme has an (average) bit rate of $1\frac{1}{2}$ bits/T. This three-level scheme is quite similar to a partial-response scheme, but since any level is available for use during any signaling interval (i.e., different transmitted symbols are independent), the data rate is greater. If we place two such 1-D schemes in quadrature, we obtain the 2-D scheme shown in Fig. 5, which achieves an (average) bit rate of 3 bits/T. This 2-D scheme provides a shaping gain of $4/3 = 1.25$ dB, with a CER2 of $9/8 = 1.125$ and a PAR2 of 2. The overall gain (including the discretization factor γ_d) is $7/6 = 0.67$ dB.

The use of a binary prefix code will not, in general, produce an optimal nonuniform scheme unless the constellation is "matched" to the code. A constellation Ω is said to be matched to some binary prefix code if, for some $\lambda \geq 0$, a Maxwell–Boltzmann distribution with parameter λ, induces probabilities on the constellation points that are all integral powers of two. This means that for some $\lambda \geq 0$,

$$e^{-\lambda \|\boldsymbol{r}\|^2}/Z(\lambda) = 2^{-l(\boldsymbol{r})} \tag{31}$$

for all $\boldsymbol{r} \in \Omega$, where $l(\boldsymbol{r})$ is a positive integer. The matching condition (31) is trivially satisfied when $\lambda = 0$ by any constellation of size 2^l. However, for positive λ, the matching condition is strong, and we expect relatively few constellations to satisfy it.

Note that, in Fig. 5, each constellation point conveys either two, three or four bits (with the outer points conveying more bits than the inner points). In particular, each point conveys at least two bits. This implies that we may consider this scheme to consist of a fixed-rate "primary" channel, conveying two bits per symbol, and a variable-rate "secondary" channel, conveying an average of one bit per symbol. In general, a binary prefix code with codeword lengths $\{l_1 \leq l_2 \leq \cdots \leq l_M\}$ assigned to an ND constellation of size M will produce a signaling scheme with an overall normalized average bit rate $\beta = (2/N)\sum_{i=1}^{M} l_i 2^{-l_i}$. The fixed primary channel rate is $\beta_p = 2l_1/N$, while the variable secondary channel rate is $\beta_s = \beta - \beta_p$. It is quite possible to have $\beta_s > \beta_p$, so the names primary and secondary do not necessarily refer to relative bit rates. In most practical circumstances, we will select

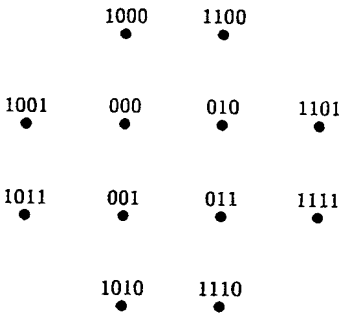

Fig. 6. A nonuniform QAM scheme. $\beta = 3.5$.

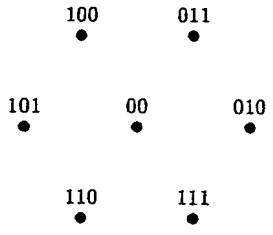

Fig. 7. A nonuniform QAM scheme. $\beta = 2.75$.

$\beta_s < \beta_p$. Note also that since the primary channel operates at a fixed rate, it can operate as a "standard" channel, and is not affected by the system problems associated with variable rate transmission. The "opportunistic secondary channels" of [2] and the "in-band coding method" of [21] (see also [17]–[20]) are examples of nonuniform signaling schemes that use prefix codes to separate data into primary and secondary channels, although these schemes are not described in terms of prefix codes.

Other examples of 2-D constellations matched to prefix codes are shown in Figs. 6 and 7. The scheme of Fig. 6 was used by Forney and Wei [2, Fig. 7(b)] to illustrate the notion of an opportunistic secondary channel, while the scheme of Fig. 7 is based on 7 points of lowest energy in the 2-D hexagonal lattice A_2. A limited search turned up several additional examples in higher dimensions. For example, the constellation containing the origin and the first shell of the lattice D_4 in four dimensions has theta series $1 + 24x$ and is matched to a binary prefix code having one codeword with two bits and 24 codewords with five bits.

B. Huffman Codes

Although, in general, a complex binary prefix code will not match a constellation in the sense of (31), we nevertheless expect prefix codes to provide shaping gain. Given an ND constellation Ω of size $|\Omega|$, in which the ith point $\boldsymbol{r}_i$ has energy $\|\boldsymbol{r}_i\|^2$, $1 \leq i \leq |\Omega|$, the optimal (gain-maximizing) complete binary prefix code with codeword lengths l_i, $1 \leq i \leq |\Omega|$, would maximize the quantity $f = (2^\beta - 1)/E$, where $\beta = (2/N)\sum_i l_i 2^{-l_i}$ and $E = (2/N)\sum_i \|\boldsymbol{r}_i\|^2 2^{-l_i}$, subject to the constraint (30). Unfortunately, short of searching all complete binary prefix codes with $|\Omega|$ codewords, we know of no general method for finding the optimal prefix code.

Rather than attempting to find the optimal code, we have taken the approach of finding approximations to the optimal

Maxwell–Boltzmann distribution with distributions in which all probabilities are positive integer powers of 1/2. (Stubley and Blake consider a more general matching problem in [34].) We refer to such approximations as "dyadic approximations" to the Maxwell–Boltzmann distribution. To find the "best" dyadic approximation, we need a measure of distance between the Maxwell–Boltzmann distribution and its dyadic approximation. A commonly used measure of distance between two probability distributions P (with probability masses $\{p_1, \cdots, p_M\}$), and Q (with probability masses $\{q_1, \cdots, q_M\}$) is the relative entropy of P with respect to Q:

$$D(P, Q) \triangleq \sum_{i=1}^{M} p_i \log_2 \frac{p_i}{q_i}.$$

When Q is dyadic, so that $q_i = 2^{-l_i}$,

$$D(P, Q) = \sum_{i=1}^{M} l_i p_i - H(P) \tag{32}$$

where $H(P) \triangleq -\sum_i p_i \log_2 p_i$ is the entropy of P.

From the point of view of source coding, $D(P, Q)$ represents the redundancy of a source code used to represent the output of a discrete memoryless source with alphabet of size M and distribution P. As is well known, the redundancy (32) is minimized by the Huffman procedure [35]. Furthermore, the Huffman procedure always results in a complete prefix code. The existence of an algorithm for minimizing $D(P, Q)$ is our primary motivation for choosing this particular measure; indeed, other measures may be more naturally suited to the problem. Nevertheless, as will become evident, this approach of minimizing $D(P, Q)$ leads to excellent gain values that can be made to approach the ultimate shaping gain.

To illustrate their performance, we have computed dyadic approximations to the Maxwell–Boltzmann distribution for two constellations based on $\mathbf{Z}^2$. The two constellations were chosen quite arbitrarily: one consists of the 21 points of least energy in $\mathbf{Z}^2$; the other consists of the 121 points of least energy. The results are illustrated in Fig. 8, and were obtained by varying the Maxwell–Boltzmann parameter λ from zero through positive values. As expected, the dyadic approximations (marked with a triangle for the 121-point constellation and a square for the 21-point constellation) have lower gain values than those obtained from the optimal Maxwell–Boltzmann distribution (the solid curves); however, the gain values do follow the general trends obtained for the Maxwell–Boltzmann distribution. Due to the flexibility afforded by having a larger number points, the larger constellation has a greater number of different dyadic approximations to the Maxwell–Boltzmann distribution.

The fact that these dyadic approximations follow the same general trends obtained for the Maxwell–Boltzmann distribution suggests the following algorithm for designing prefix codes to achieve shaping gain. Given an N-D constellation, we proceed as follows.

1) From the constellation theta series, we numerically determine the value of the Maxwell–Boltzmann parameter λ that maximizes gain G (8). Call this value λ_{opt}.

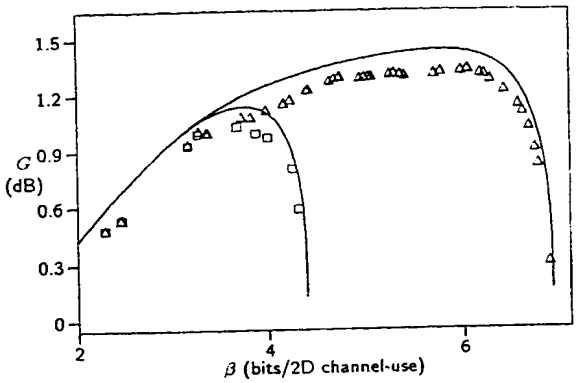

Fig. 8. Illustrating the performance of dyadic approximations to the Maxwell–Boltzmann distribution obtained from the Huffman procedure.

2) Using λ_{opt}, we generate a list of Maxwell–Boltzmann probabilities

$$p_i = p(\boldsymbol{r}_i) = \exp(-\lambda_{\text{opt}} \|\boldsymbol{r}_i\|^2)/Z(\lambda_{\text{opt}})$$

for all $\boldsymbol{r}_i \in \Omega$.

3) We apply the Huffman procedure to the list p_i to obtain a complete binary prefix code. The performance of this code is then evaluated.

Note that, in general, the Huffman procedure does not result in a unique code. We have chosen the version of the Huffman procedure, described in [36, p. 68], that results in least variation among the codeword lengths. Note also that we have chosen λ_{opt} to maximize total gain. Since the coding gain is fixed for a given lattice, this is equivalent to maximizing the product $\gamma_s\gamma_d$. It is important to note that this is *not* equivalent to maximizing the shaping gain γ_s since this, in principle, can be accomplished by making λ arbitrarily large.

We have applied this procedure to spherical constellations based on the 2-D $\mathbf{Z}^2$ lattice and its translate $(\mathbf{Z} + 1/2)^2$, as well as the 2-D hexagonal lattice A_2. The results are given in Tables I–III. Each constellation Ω consists of the $|\Omega|$ points of least energy drawn from the corresponding infinite lattice (or translate). Shown in each table are the primary and secondary bit rates (β_p and β_s) obtained from the Huffman code. Each table lists the parameter $\gamma_s\gamma_d$, obtained by dividing the total gain of the signaling scheme by the coding gain of the lattice upon which it is based. The 2-D peak-to-average energy ratio PAR2 and the 2-D constellation expansion ratio CER2 are listed. In addition, the parameter N_{eff}, the "effective dimension," is listed. We define the effective dimension of a shaping scheme to be the smallest dimension N for which the shaping gain of an N-sphere $\gamma_\otimes(N)$ (39) (properly multiplied by γ_d) meets or exceeds the shaping gain provided by the scheme in question, i.e.,

$$N_{\text{eff}}(\gamma_s, \beta) \triangleq \min\{N \in \mathbf{Z}: \gamma_\otimes(N)(1 - 2^{-\beta}) \geq \gamma_s\}.$$

As can be seen from the tables, very satisfactory shaping gain values, with effective dimensions numbering in the hundreds of dimensions, are obtained from these Huffman prefix codes.

The shaping gains obtained from these Huffman codes seem to be the highest ever reported, exceeding those reported in [8, Table IV]. As previously noted, however, maximization

TABLE I
PERFORMANCE OF HUFFMAN-CODED SIGNAL CONSTELLATIONS BASED ON $\mathbb{Z}^2$

$\lvert\Omega\rvert$	β_p	β_s	$\gamma_s\gamma_d$ (dB)	N_{eff}	PAR2	CER2
9	2	1.000	0.669	47	2.000	1.125
13	2	1.125	0.834	95	3.765	1.490
21	3	0.688	1.030	111	3.200	1.630
25	3	0.906	1.039	80	4.357	1.667
29	3	1.078	1.069	78	4.347	1.717
37	3	1.359	1.172	120	4.025	1.803
45	3	1.641	1.39	159	4.333	1.804
49	3	1.656	1.249	173	5.285	1.943
57	3	1.719	1.263	189	5.386	2.165
61	4	0.922	1.259	139	4.923	2.012
69	4	1.020	1.275	149	5.120	2.127
81	4	1.281	1.289	135	5.327	2.083
89	4	1.291	1.301	152	5.517	2.273
97	4	1.389	1.288	124	5.724	2.315
101	4	1.667	1.285	103	5.182	1.988
109	4	1.708	1.301	115	5.368	2.085
113	4	1.710	1.303	117	5.679	2.159
121	4	1.780	1.318	129	5.573	2.202
129	4	1.784	1.323	136	6.015	2.341
137	5	1.007	1.340	145	5.291	2.131
145	5	1.077	1.359	171	5.550	2.148
149	5	1.078	1.360	174	6.042	2.205
161	5	1.119	1.367	184	5.998	2.316
169	5	1.168	1.369	184	6.031	2.350
177	5	1.256	1.373	186	5.784	2.316
185	5	1.304	1.377	190	6.124	2.341
193	5	1.319	1.379	194	6.376	2.417
197	5	1.336	1.379	192	6.611	2.438

TABLE II
PERFORMANCE OF HUFFMAN-CODED SIGNAL CONSTELLATIONS BASED ON $(\mathbb{Z} + 1/2)^2$

$\lvert\Omega\rvert$	β_p	β_s	$\gamma_s\gamma_d$ (dB)	N_{eff}	PAR2	CER2
12	2	1.000	0.670	47	2.500	1.500
16	2	1.375	0.757	37	3.429	1.542
24	3	0.844	1.045	93	3.714	1.672
32	3	1.094	1.136	130	4.121	1.874
44	4	0.734	1.123	58	3.791	1.653
52	4	0.820	1.190	85	4.199	1.841
60	4	0.969	1.232	104	4.862	1.916
68	4	1.086	1.253	113	4.979	2.002
76	4	1.266	1.289	138	4.848	1.976
80	4	1.395	1.320	172	4.851	1.902
88	4	1.412	1.334	202	5.198	2.067
96	4	1.432	1.341	219	5.911	2.224
112	4	1.691	1.340	175	5.239	2.167
120	4	1.828	1.353	185	5.358	2.112
124	5	0.945	1.376	232	5.502	2.012
140	5	0.955	1.385	268	5.747	2.257
148	5	1.083	1.397	293	5.514	2.183
156	5	1.157	1.406	328	5.716	2.186
164	5	1.166	1.410	351	5.921	2.284
172	5	1.223	1.414	364	6.146	2.303
180	5	1.224	1.417	384	6.369	2.408
188	5	1.232	1.418	390	6.558	2.500
192	5	1.245	1.416	373	6.719	2.531
208	5	1.378	1.398	247	6.297	2.501
216	5	1.449	1.389	209	6.554	2.473
232	5	1.511	1.385	193	6.633	2.543
240	5	1.654	1.389	192	6.172	2.383
248	5	1.726	1.398	209	6.041	2.343
256	5	1.733	1.400	214	6.171	2.407

of shaping gain γ_s itself is not our aim; rather, we have attempted to maximize the combination $\gamma_s\gamma_d$. Further numerical calculations obtained from dyadic approximations to the Maxwell–Boltzmann distribution with parameter $\lambda > \lambda_{\text{opt}}$ show that, in some cases, the effective dimension N_{eff} can be made to increase significantly above the values shown in Tables I–III. Note also that PAR2 and CER2 values shown in Tables I–III are all quite reasonable, especially when compared to the PAR2 and CER2 of large-dimensional Voronoi constellations [6]. The PAR2 and CER2 can, in principle, be improved by sacrificing some shaping gain. Indeed, numerical calculations show that dyadic approximations to the Maxwell–Boltzmann distribution with parameter $\lambda < \lambda_{\text{opt}}$ will, in general, result in improved PAR2 and CER2, with some corresponding sacrifice in overall gain. We have also applied this procedure to multidimensional constellations, with similar results. However, since many multidimensional constellations are best implemented as coset codes (see [10] and [11]), it may be preferable to use a nonuniform signal point selection scheme that is suited for a coset code (as described in the next section) rather than a direct mapping of the words of a prefix code onto the constellation points.

It follows from standard arguments in information theory (e.g., [1, Section 5.4]) that the redundancy of the optimal code for a discrete memoryless source can be made to approach zero by considering Cartesian products of the source. This implies that the relative entropy between the Maxwell–Boltzmann distribution and its dyadic approximation can be made to approach zero by considering Cartesian products of the basic constellation. Convergence in relative entropy implies L_1 convergence of the probabilities [1, Section 12.6]; hence, the performance obtained from our dyadic approximations can be made to approach arbitrarily closely to the performance obtained by using the optimal Maxwell–Boltzmann distribution. Numerical calculations confirm the performance improvement obtained by working with Cartesian products of the basic constellations.

VIII. CODED NONUNIFORM SIGNALING

In this section we study how optimal nonuniform signaling fits into the general framework of coset codes, first introduced by Calderbank and Sloane [12] and extensively studied by Forney [10], [11].

A. Memoryless Signal Point Selectors

A signaling scheme based on a coset code has two components as shown in Fig. 1. A coset code $\mathbb{C}$, based on the

TABLE III
PERFORMANCE OF HUFFMAN-CODED SIGNAL CONSTELLATIONS BASED ON A_2

$\lvert\Omega\rvert$	β_p	β_s	$\gamma_s\gamma_d$ (dB)	$\mathcal{N}_{\text{eff}}$	PAR2	CER2
7	2	0.750	0.423	27	1.333	1.041
13	2	1.031	0.732	62	3.429	1.590
19	2	1.453	0.884	62	3.413	1.735
31	3	1.187	1.047	59	3.584	1.701
37	3	1.477	1.142	81	3.815	1.662
43	3	1.668	1.242	157	4.531	1.691
55	3	1.754	1.292	268	4.668	2.038
61	4	1.006	1.304	214	4.808	1.899
73	4	1.049	1.324	273	5.562	2.205
85	4	1.449	1.286	117	4.583	1.946
91	4	1.473	1.296	126	5.376	2.049
97	4	1.557	1.314	143	5.496	2.061
109	4	1.695	1.330	155	5.186	2.104
121	4	1.766	1.337	161	5.472	2.224
127	5	0.851	1.344	163	5.992	2.201
139	5	0.912	1.351	171	5.907	2.308
151	5	1.078	1.368	191	5.564	2.236
163	5	1.159	1.379	211	5.809	1.282
169	5	1.167	1.380	214	6.448	2.352
187	5	1.280	1.379	199	6.082	2.407
199	5	1.441	1.386	200	5.772	2.291
211	5	1.487	1.393	215	6.134	2.352
223	5	1.526	1.396	222	6.395	2.420

partitioning of a lattice Λ (possibly translated by some constant vector) into the cosets of sublattice Λ', produces a sequence of sets of channel symbols, drawn from the alphabet of the cosets of Λ' in Λ. The actual transmitted constellation point is determined by the signal point selector $\mathbf{S}$. As discussed in Section II, both the coset code $\mathbf{C}$ and the signal point selector $\mathbf{S}$ contribute to the transmission of data. It is important to note that, as nonuniform signaling is a shaping technique, for the schemes we propose only the signal point selector $\mathbf{S}$ is affected. The coset code $\mathbf{C}$ is unchanged relative to well-known schemes such as those of Ungerboeck [37].

The simplest type of signal point selector is memoryless, or time invariant. When $\mathbf{S}$ is memoryless, each time a coset of Λ' is made available to $\mathbf{S}$, the subset from which the constellation point is selected is the same, and the choice is made independently. For example, the signal point selector could always choose from the K points of least norm in each coset. For a block coset code $\mathbf{C}$ based on a 2-D lattice Λ, this would result in a polydisc-shaped constellation. Cubic constellations are achieved if $\mathbf{S}$ always selects from a square-shaped region in each coset.

More complicated signal point selectors have memory, i.e., they are time-varying. To achieve generalized cross constellations [2], Voronoi constellations [6], or indeed constellations based on any region that is not a Cartesian product of lower-dimensional regions, the signal point selector must be time-varying. The block shaping codes of Calderbank and Ozarow [7] and the trellis shaping codes of Forney [8] are examples of time-varying signal point selectors. For coset codes based on 2-D lattices, and assuming uniform signaling, time-varying signal point selectors are necessary to achieve shaping gains that exceed the shaping gain of a 2-D disc. Indeed, the best possible shaping gain in N-space is achieved by an N-sphere, a region not decomposable as a Cartesian product of lower-dimensional regions.

As discussed in Section VI, under nonuniform signaling, the best regions with which to shape a constellation are polydiscs, as these achieve a given shaping gain with least shaping constellation expansion ratio CER2$_s$ and least peak-to-average energy ratio PAR2. Since polydiscs are by definition a product of 2-D discs, polydisc-shaped constellations can be achieved by a memoryless signal point selector combined with a coset code based on a 2-D lattice. We focus our attention, therefore, on memoryless nonuniform signal point selectors.

B. Coset Codes

The coset codes considered in this paper are based on any L-way partition of Λ/Λ' of an ND lattice Λ into the L cosets of a sublattice Λ'. We focus our attention on binary coset codes, where $L = 2^{k+r}$, although generalization to nonbinary coset codes is straightforward.

Let C be a binary rate-$k/(k+r)$ encoder that takes in k bits per ND and puts out $k+r$ coded bits. These coded bits can be used to select one of the 2^{k+r} cosets of Λ' in Λ. The resulting coset code is denoted $\mathbf{C}(\Lambda/\Lambda'; C)$. When the binary encoder C is a block code, the coset code $\mathbf{C}(\Lambda/\Lambda'; C)$ defines a finite-dimensional sphere packing; often this sphere packing is actually a lattice. When C is a convolutional code, the resulting coset code is a trellis code.

Let us denote a set of $L = 2^{k+r} = |\Lambda/\Lambda'|$ coset leaders of the cosets of Λ' in Λ by $\{\boldsymbol{c}_1, \boldsymbol{c}_2, \cdots, \boldsymbol{c}_L\}$. Each time the memoryless signal point selector $\mathbf{S}$ is presented with ith coset $\boldsymbol{c}_i + \Lambda'$, it selects some point for transmission. The set of all possible points drawn from the ith coset forms the ith constellation Ω_i.

Given that $\mathbf{S}$ is presented with the ith coset, if a point $\boldsymbol{r}_i \in \boldsymbol{c}_i + \Lambda'$ is selected with probability $p(\boldsymbol{r}_i)$, then we can determine a normalized average bit rate $\beta_i = -2\sum_{\boldsymbol{r}_i \in \Omega_i} p(\boldsymbol{r}_i) \log_2 [p(\boldsymbol{r}_i)]/N$ and a normalized average energy $E_i = 2\sum_{\boldsymbol{r}_i \in \Omega_i} p(\boldsymbol{r}_i)\|\boldsymbol{r}_i\|^2/N$ for the ith constellation. If the coset code $\mathbf{C}$ selects the ith coset with probability $P[i]$ (and usually $P[i] = 1/L$), then the normalized average number of bits taken in by the signal point selector $\mathbf{S}$ is $\kappa(\mathbf{S}) = \sum_{i=1}^{L} P[i]\beta_i$, and the normalized average energy is $E = \sum_{i=1}^{L} P[i]E_i$. Since the code $\mathbf{C}$ takes in k bits per N dimensions, or $\kappa(\mathbf{C}) = 2k/N$ bits per two dimensions, the overall normalized bit rate $\beta = \kappa(\mathbf{C}) + \kappa(\mathbf{S})$. If the coset code $\mathbf{C}$ has minimum squared Euclidean distance $d^2_{\min}$, then, from (8) the gain G of the coded modulation scheme may be written as

$$G = \left(\frac{2^{\kappa(\mathbf{C})}\, d^2_{\min}}{V(\Lambda')^{2/N}}\right)\left(\frac{2^{\kappa(\mathbf{S})} V(\Lambda')^{2/N}}{6E}\right)(1 - 2^{-\beta})$$
$$= \gamma_c(\mathbf{C})\gamma_s(\mathbf{S})\gamma_d(\beta).$$

Here, as in (9), we have separated the total gain into the product of the coding gain $\gamma_c(\mathbf{C})$ of the coset code $\mathbf{C}$ [10],

the overall shaping gain $\gamma_s(\mathbf{S})$ provided by the signal point selector $\mathbf{S}$, and a discretization factor $\gamma_d(\beta)$.

C. Continuous Approximations

Suppose now that each constellation Ω_i is obtained from the intersection of the ith coset of Λ' with the same finite region $\mathbf{R}$. Further, suppose that the memoryless signal point selector $\mathbf{S}$ selects each point in Ω_i with a Maxwell–Boltzmann distribution, i.e., given that $\mathbf{S}$ is presented with the ith constellation, a point $\mathbf{r}_i \in \Omega_i$ is selected with probability $p(\mathbf{r}_i) = \exp(-\lambda\|\mathbf{r}_i\|^2)/Z_i(\lambda)$, where λ is fixed for all constellations.

Using the same continuous approximation principles as in Section VI, the average energy E_i and bit rate β_i for each subconstellation can be estimated via (21) and (23) from a continuous Gaussian distribution, truncated to the region $\mathbf{R}$. It follows that each subconstellation supports approximately the same bit rate, at the same cost in average energy. Thus, independently of the coset code $\mathbf{C}$ and the probability with which each coset is selected, $\kappa(\mathbf{S}) \approx \beta_i(\mathbf{R}, \lambda)$ and $E \approx E_i(\mathbf{R}, \lambda)$.

Using these estimates for bit rate and average energy, the shaping gain $\gamma_s(\mathbf{S})$ is estimated by

$$\gamma_s(\mathbf{R}, \lambda) = \frac{2^{2H(\mathbf{R}, \lambda)/N}}{6E(\mathbf{R}, \lambda)},$$

which is the same expression as (25). Similarly, we find that CER2, the 2-D constellation expansion ratio, is approximately the produce of a coding constellation expansion ratio $\mathrm{CER2}_c(\mathbf{C})$ and a shaping constellation expansion ratio $\mathrm{CER2}_s(\mathbf{S})$. In terms of the normalized redundancy $\rho(\mathbf{C}) \triangleq 2r/N$ of the binary encoder $\mathbf{C}$, we have

$$\mathrm{CER2}_c(\mathbf{C}) = 2^{\rho(\mathbf{C})}\frac{V(\Lambda)^{2/N}}{V(\Lambda_2)},$$

while

$$\mathrm{CER2}_s(\mathbf{S}) = 2^{\rho(\mathbf{R}, \lambda)}\frac{V(\mathbf{R}_2)}{V(\mathbf{R})^{2/N}},$$

exactly as in Section VI. Approximations for the 2-D peak-to-average energy ratio PAR2 also lead to expressions identical to those given in Section VI.

In general, to the accuracy of the continuous approximation, the shaping gain $\gamma_s(\mathbf{R}, \lambda)$ achieved by our memoryless nonuniform signal point selector $\mathbf{S}$ is completely independent of the choice of coset code $\mathbf{C}$. Trade-offs involving shaping constellation expansion ratio $\mathrm{CER2}_s(\mathbf{S})$ and PAR2 are also completely independent of the choice of coset code $\mathbf{C}$. As noted, it is desirable to have constellations shaped as polydiscs, since these achieve a given shaping gain with minimum $\mathrm{CER2}_s$ and minimum PAR2. Therefore, a nonuniform signaling scheme based on a multidimensional lattice Λ is perhaps best implemented as a coset code involving the constituent 2-D lattice Λ_2, in which the ith constellation Ω_i is circular and the signal point selector $\mathbf{S}$ chooses from Ω_i according to a Maxwell–Boltzmann distribution.

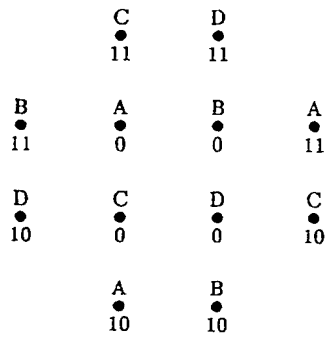

Fig. 9. The sequence of subconstellations {A, B, C, D} is determined by an Ungerboeck code. Signal point selection within each subconstellation is performed using a simplex prefix code.

D. Memoryless Signal Point Selection with Huffman Codes

As in the case of uncoded transmission, probably the simplest method of selecting points from the coset sequences generated by a coset code is through a complete binary prefix code. Because we would like our constellations to be polydiscs, and we have restricted our attention to binary coset codes, we focus on coset codes based on the 2-D lattice $\mathbf{Z}^2$.

A simple example of combining a nonuniform memoryless signal point selector with a trellis code is shown in Fig. 9. The trellis code is a simple four-state Ungerboeck code [37] based on a translate of the four-way partition $\mathbf{Z}^2/2\mathbf{Z}^2$; this code provides a coding gain $\gamma_c = 2 = 3.01$ dB. Each subconstellation (labeled A, B, C, or D) consists of a single inner point at squared Euclidean distance 1/2 from the origin and two outer points at squared Euclidean distance 5/2 from the origin. By using the prefix code $\{0, 10, 11\}$, the signal point selector chooses the inner point with probability 1/2 and each outer point with probability 1/4. The transmitted rate β for this scheme is 2.5 bits with a primary channel rate of 2 bits and a secondary channel rate of 0.5 bits. The shaping gain $\gamma_s = 0.994$ dB. The overall gain $G = \gamma_c\gamma_s\gamma_d = 3.16$ dB.

The four-way partition $\mathbf{Z}^2/2\mathbf{Z}^2$, translated as in Fig. 9, yields four cosets with identical weight distributions. The weight enumerator for these cosets begins

$$\Theta_\Omega(x) = x^{1/2} + 2x^{5/2} + x^{9/2} + 2x^{13/2} + 2x^{17/2} + \cdots.$$

As in Section VII, we have computed dyadic approximations to the optimal Maxwell–Boltzmann distribution by applying the Huffman procedure. The results are given in Table IV. We have assumed a coset code $\mathbf{C}(\mathbf{Z}^2/2\mathbf{Z}^2; C)$ with a normalized bit rate $\kappa(\mathbf{C}) = 1$. This bit rate is included in the primary rate β_p of Table IV. The overall constellation, of size $|\Omega|$, is the union of four subconstellations, each of size $|\Omega|/4$. As in Tables I–III, the gain $\gamma_s\gamma_d$ and the effective dimension N_{eff} are listed, along with PAR2 and $\mathrm{CER2}_s$. Note that $\mathrm{CER2}_s(\mathbf{C}) = 2$. As before, very satisfactory shaping gains are achievable via these prefix codes, with N_{eff} numbering in the hundreds of dimensions for the larger constellations.

Although the prefix codes for the smaller constellations have smaller N_{eff}, as discussed at the end of Section VII, applying the Huffman procedure to Cartesian products of these subconstellations can improve the gain. Note, however, that the

TABLE IV
PERFORMANCE OF HUFFMAN-CODED RATE 1/2 COSET CODES BASED ON $\mathbf{Z}^2/2\mathbf{Z}^2$

$\|\Omega\|$	β_p	β_s	$\gamma_s\gamma_d$ (dB)	$\bar{N}_{\text{eff}}$	PAR 2	CER2$_s$
12	2	0.500	0.149	17	1.667	1.061
16	2	0.750	0.378	23	2.571	1.189
24	3	0.375	0.555	18	2.364	1.157
32	3	0.562	0.798	34	2.833	1.354
44	3	0.906	0.948	46	3.333	1.467
52	3	1.281	1.064	60	3.013	1.337
60	3	1.328	1.128	87	3.771	1.494
68	3	1.375	1.168	112	4.075	1.639
76	3	1.562	1.155	82	3.892	1.608
80	3	1.570	1.174	94	4.232	1.684
88	4	0.883	1.221	102	3.695	1.491
96	4	0.906	1.247	126	4.207	1.601
112	4	0.938	1.289	192	4.426	1.827
120	4	1.180	1.295	158	4.186	1.655
124	4	1.184	1.303	173	4.641	1.706
140	4	1.207	1.329	239	4.818	1.895
148	4	1.213	1.340	280	5.036	1.995
156	4	1.432	1.292	125	4.647	1.807
164	4	1.455	1.303	137	4.770	1.869
172	4	1.494	1.306	137	5.011	1.908
180	4	1.656	1.308	127	4.634	1.785
188	4	1.705	1.321	140	4.649	1.802
192	4	1.707	1.326	147	4.807	1.838
208	4	1.717	1.342	174	4.950	1.977
216	4	1.728	1.345	179	5.387	2.038
232	5	0.924	1.342	153	4.961	1.911
240	5	0.941	1.348	162	5.042	1.953
248	5	0.986	1.353	169	5.023	1.956
256	5	1.020	1.357	174	5.040	1.973

signal point selector $\mathbf{S}$ is no longer memoryless in this case. Other partitions of $\mathbf{Z}^2$, e.g., the 8-way partition $\mathbf{Z}^2/2R\mathbf{Z}^2$ or the 16-way partition $\mathbf{Z}^2/4\mathbf{Z}^2$, allow the use of more powerful coset codes, many of which are listed in [10, Tables IV, V, IX, X, XI]. Unlike the 4-way partition $\mathbf{Z}^2/2\mathbf{Z}^2$, the cosets in these partitions do not all have the same theta series. For example, the 8-way partition of $\mathbf{Z}^2 + (1/2, 1/2)$, has two classes of cosets [typified by $A = 2R\mathbf{Z}^2 + (1/2, 1/2)$ and $B = 2R\mathbf{Z}^2 + (3/2, 1/2)$] with weight enumerators

$$\Theta_A(x) = x^{1/2} + x^{9/2} + 2x^{17/2} + 3x^{25/2} + \cdots$$
$$\Theta_B(x) = 2x^{5/2} + 2x^{13/2} + 2x^{29/2} + 2x^{37/2} + \cdots,$$

respectively. Each class consists of four different cosets. Similarly, the 16-way partition of $\mathbf{Z}^2 + (1/2, 1/2)$ has three classes of cosets [typified by $A = 4\mathbf{Z}^2 + (1/2, 1/2)$, $B = 4\mathbf{Z}^2 + (3/2, 1/2)$, and $C = 4\mathbf{Z}^2 + (3/2, 3/2)$] with weight enumerators

$$\Theta_A(x) = x^{1/2} + 2x^{25/2} + 2x^{41/2} + x^{49/2} + \cdots$$
$$\Theta_B(x) = x^{5/2} + x^{13/2} + x^{29/2} + x^{37/2} + \cdots$$
$$\Theta_C(x) = x^{9/2} + 2x^{17/2} + x^{25/2} + 2x^{65/2} + \cdots,$$

respectively. Class A consists of four cosets, class B consists of eight cosets, and class C consists of four cosets. We have found good dyadic approximations to the Maxwell–Boltzmann distribution for these partitions, with results similar to those given in Table IV. Many different schemes may be designed by combining the various signal point selectors obtained. For example, it may be advantageous for different cosets to convey different numbers of secondary channel bits.

In general, by applying the Huffman algorithm to obtain dyadic approximations to Maxwell–Boltzmann distribution, we have been able to obtain a variety of different schemes. By varying λ, schemes that trade off gain for improved PAR2 and CER2$_s$ are easily obtained.

IX. DISCUSSION AND CONCLUSION

From the point of view of coded modulation, we have seen that nonuniform signal point selection is an energy-minimizing or shaping operation. When constellation points are selected with Maxwell–Boltzmann probabilities, the ultimate in shaping gain performance can be achieved in *any* dimension. Dyadic approximations to the optimal Maxwell–Boltzmann distribution are easily obtained by applying the Huffman procedure. The performance of the resulting shaping schemes is often close to optimum, with effective dimensions numbering in the hundreds, and can be made to approach the optimum by considering Cartesian products of the basic constellations. By varying the Maxwell–Boltzmann parameter, trade-offs between shaping gain, 2-D constellation expansion ratio or 2-D peak-to-average energy ratio are easily accomplished. In a sense, the implementation complexity for these schemes is trivial, since data to constellation point mappings (and vice versa) are easily performed by table lookup. Furthermore, these schemes are easily incorporated into well-known lattice-type coded modulation schemes. All of these properties make nonuniform signaling very attractive.

The principal drawback, as pointed out at the outset, is the variable bit rate. While only the secondary channel data are subject to the problems associated with buffer under- and overflow and the insertion and deletion of bits in the decoded bit stream, these problems may be acceptable only in certain applications, e.g., for the transmission of internal control signals. Left unsolved, these problems will tend to limit the broad applicability of nonuniform signaling. Solving the system problems associated with nonuniform signaling will certainly increase the complexity of implementation. Yet, in order to achieve the large shaping gains achieved with nonuniform signaling, uniform signaling schemes will themselves tend to become quite complex (see [8, Table IV]). It remains an open problem to evaluate and compare these complexities.

APPENDIX A
ERROR COEFFICIENTS FOR BASELINE CONSTELLATIONS

In this appendix, we compute the average nearest neighbor multiplicity (error coefficient) for a cubic constellation of side M drawn from the lattice $\mathbf{Z}^N$. This problem is easily solved using the notion of a nearest neighbor enumerator.

Given a finite constellation Ω, recall that $N_{\min}(\boldsymbol{r})$ denotes the number of points of Ω at distance $d_{\min}$ from the point $\boldsymbol{r}$.

Define $A(x) \triangleq \sum_{r\in\Omega} x^{N \min(r)}$; then $A(x)$ is a polynomial with integer coefficients that we call the nearest neighbor enumerator for Ω. For example, a simple 1-D PAM constellation with M points has $A(x) = 2x + (M-2)x^2$, indicating that two points of the constellation have a single nearest neighbor, while $M-2$ points have two nearest neighbors.

Assuming uniform signal point selection, it is easy to see that average nearest neighbor multiplicity $\overline{N}$ is given by $A'(1)/A(1)$ where $A'(x) = dA(x)/dx$. Thus for the M-PAM example of the previous paragraph, $\overline{N} = 2(1-1/M)$.

It is easily seen that the nearest neighbor enumerator for Ω^n, the n-fold Cartesian product of Ω with itself, is $A(x)^n$. The average nearest neighbor multiplicity of the Cartesian product is then $nA(1)^{n-1}A'(1)/A(1)^n = nA'(1)/A(1)$; in other words, taking the n-fold Cartesian product of a constellation with itself multiplies the average nearest neighbor multiplicity by n. Since an N-D cubic constellation of side M is the N-fold Cartesian product of simple 1-D PAM constellations, we have $\overline{N} = 2N(1-1M)$ for such constellations. Furthermore, since $M = 2^{\beta/2}$, we obtain $\overline{N}_\oplus = 2N(1-2^{-\beta/2})$ for our N-D baseline constellations.

Appendix B
Continuous Approximations for Spherical Constellations

In this appendix, we specialize the continuous approximations derived in Section VI to the case where $\mathbf{R} = \mathbf{B}_N(R)$, an N-ball of radius R centered at the origin. Cubic constellations can be considered to be Cartesian products of 1D "spherical constellations" and so are (by the separability of Gaussian densities) a special case. By letting $R \to \infty$, we obtain continuous approximations for the case of an infinite constellation.

Many of the expressions derived in this Appendix may be written in terms of the (normalized) incomplete Gamma function $P(a, x)$, defined in [23] as

$$P(a, x) \triangleq \frac{1}{\Gamma(a)}\int_0^x t^{a-1}e^{-t}\,dt,$$

where we note that $\lim_{x\to\infty} P(a, x) = 1$.

For a finite N-D spherical constellation Ω, in our estimates we choose the spherical radius R so that $V(\mathbf{B}_N(R)) = |\Omega|V(\Lambda)$, i.e., so that approximation (19) holds with equality.

Energy: From (21) we obtain

$$E(\mathbf{B}_N(R), \lambda) = \frac{1}{\lambda}\cdot\frac{P(N/2+1, \lambda R^2)}{P(N/2, \lambda R^2)}. \tag{33}$$

Since $E(\mathbf{B}_N(R), 0) = 2R^2/(N+2)$, we find that the energy savings factor (29) is

$$g_E(\mathbf{B}_N(R), \lambda) = \frac{2\lambda R^2 P(N/2, \lambda R^2)}{(N+2)P(N/2+1, \lambda R^2)}. \tag{34}$$

Bit Rate: The entropy of a continuous Gaussian random variable (with parameter λ) truncated to $\mathbf{B}_N(R)$ is given by

$$H(\mathbf{B}_N(R), \lambda) = \log_2[(\pi/\lambda)^{N/2}P(N/2, \lambda R^2)] + N\lambda E(\mathbf{B}_N(R), \lambda)/(2\ln 2).$$

Setting $\lambda = 0$ yields

$$H(\mathbf{B}_N(R), 0) = \log_2[V(\mathbf{B}_N(R))] = \log_2[(\pi R^2)^{N/2}/\Gamma(N/2+1)].$$

Combining these expressions gives an estimate for the normalized redundancy (27), namely,

$$\rho(\mathbf{B}_N(R), \lambda) = \log_2\left(\frac{\lambda R^2}{[\Gamma(N/2+1)P(N/2, \lambda R^2)]^{2/N}}\right) - \frac{P(N/2+1, \lambda R^2)}{P(N/2, \lambda R^2)}\log_2 e. \tag{35}$$

The bit rate can be estimated via (23) or via (27). For a spherical constellation of size Ω, we use the approximation $\beta(\mathbf{B}_N(R), \lambda) = (2/N)\log_2|\Omega| - \rho(\mathbf{B}_N(R), \lambda)$.

Shaping and Biasing Gains: Substituting (34) and (35) into (28) gives an estimate for the biasing gain. Multiplying the biasing gain by $\gamma_\otimes(N) \triangleq \pi(N/2+1)/[6\Gamma(N/2+1)^{2/N}]$, the shaping gain of an N-sphere under uniform signaling [2], yields an estimate for the shaping gain under nonuniform signaling. Explicitly, the shaping gain $\gamma_s(\mathbf{B}_N(R), \lambda)$ is

$$\gamma_s(\mathbf{B}_N(R), \lambda) = \frac{\pi}{6}\exp\left(\frac{P(N/2+1, \lambda R^2)}{P(N/2, \lambda R^2)}\right)\cdot\frac{P(N/2, \lambda R^2)^{(2/N)+1}}{P(N/2+1, \lambda R^2)}. \tag{36}$$

PAR2: The constituent 2-D constellation of $\mathbf{B}_N(R)$ (as defined in Section III) is a 2-D disc $\mathbf{B}_2(R)$ (with peak energy R^2) when N is even, and a square $\mathbf{B}_1^2(R)$ of side $2R$ (with peak energy $2R^2$) when N is odd; this may be expressed compactly by writing $r^2_{\max}(\mathbf{B}_N(R)) = [3-(-1)^N]R^2/2$. Since the normalized average energy under uniform signaling is $2R^2/(N+2)$, we have

$$\mathrm{PAR2}(\mathbf{B}_N(R), 0) = (3-(-1)^N)(N+2)/4.$$

Under nonuniform signaling this PAR2 is increased by a factor of g_E so that

$$\mathrm{PAR2}(\mathbf{B}_N(R), \lambda) = \frac{(3-(-1)^N)\lambda R^2 P(N/2, \lambda R^2)}{2P(N/2+1, \lambda R^2)}. \tag{37}$$

CER2$_s$: Under uniform signaling, a large even-dimensional spherical constellation induces a shaping constellation expansion ratio of $\mathrm{CER2}_s = \Gamma(N/2+1)^{2/N}$ (see [2]), while a large odd-dimensional spherical constellation has $\mathrm{CER2}_s = (4/\pi)\Gamma(N/2+1)^{2/N}$. Under nonuniform signaling, this $\mathrm{CER2}_s$ is increased by a factor of $2^{\rho(\mathbf{B}_N(R)\,\lambda)}$, so that

$$\mathrm{CER2}_s(\mathbf{B}_N(R), \lambda) = \frac{\pi+4+(-1)^N(\pi-4)}{2\pi}\cdot\frac{\lambda R^2}{P(N/2, \lambda R^2)^{2/N}}\cdot\exp\left[-\frac{P(N/2+1, \lambda R^2)}{P(N/2, \lambda R^2)}\right]. \tag{38}$$

Limiting Case: When $R \to \infty$, the truncated Gaussian random variable approaches a standard (untruncated) Gaussian random variable. For large R, the normalized average energy $E(\mathbf{B}_N(R), \lambda) \to 1/\lambda$, independently of N. Similarly, the energy savings factor $g_E(\mathbf{B}_N(R), \lambda) \to 2\lambda R^2/(N+2)$, the normalized redundancy $\rho(\mathbf{B}_N(R), \lambda) \to \log_2(\lambda R^2/e\Gamma(N/2+1)^{2/N})$. Thus the biasing gain

$$\gamma_b(\mathbf{B}_N(R), \lambda) \to 2e\Gamma(N/2+1)^{2/N}/(N+2) = \pi e/[6\gamma_\otimes(N)],$$

where

$$\lambda_\otimes(N) \triangleq \pi(N/2+1)/[6\Gamma(N/2+1)^{2/N}], \qquad (39)$$

is the shaping gain of the N-sphere under uniform signaling. The shaping gain $\gamma_s(\mathbf{B}_N(R), \lambda) \to \pi e/6$, *independently* of the dimension N. Of course, the biasing gain depends on N, and approaches unity as $N \to \infty$. For large values of R, $\mathrm{CER2}(\mathbf{B}_N(R), \lambda) \to \lambda R^2/e$ for even-dimensional spherical constellations, and $\mathrm{CER2}(\mathbf{B}_N(R), \lambda) \to 4\lambda R^2/(\pi e)$ for odd-dimensional spherical constellations. Similarly, for even-dimensional spherical constellations, $\mathrm{PAR2}(\mathbf{B}_N(R), \lambda) \to \lambda R^2$, independently of N. Of course, this is to be expected because the peak energy value is approximately R^2, and the average energy is approximately $1/\lambda$. Similarly, for odd-dimensional spherical constellations, $\mathrm{PAR2}(\mathbf{B}_N(R), \lambda) \to 2\lambda R^2$ because the peak energy value (in two dimensions) is approximately $2R^2$.

ACKNOWLEDGMENT

We are grateful to G. D. Forney, Jr., for his detailed comments on earlier versions of this paper, for his insightful suggestions, and for providing us with some related reference material. Comments made by the anonymous reviewers proved most helpful as well. We would also like to thank P. R. Stubley for pointing out that the Huffman algorithm minimizes the relative entropy between a distribution and its dyadic approximation.

REFERENCES

[1] T. M. Cover and J. A. Thomas, *Elements of Information Theory.* New York, Wiley, 1991.
[2] G. D. Forney, Jr., and L.-F. Wei, "Multidimensional constellations—Part I: Introduction, figures of merit, and generalized cross constellations, *IEEE J. Select. Areas Commun.*, vol. 7, pp. 877–892, Aug. 1989.
[3] E. Schrödinger, *Statistical Thermodynamics.* Cambridge: Cambridge University Press, 1962.
[4] R. K. Pathria, *Statistical Mechanics.* Elmsford, NY: Pergamon, 1972.
[5] D. Ruelle, "Thermodynamic formalism," in *Encyclopedia of Mathematics and Its Applications, vol. 5.* Reading, MA: Addison-Wesley, 1978.
[6] G. D. Forney, Jr., "Multidimensional constellations—Part II: Voronoi constellations," *IEEE J. Select. Areas Commun.*, vol. 7, pp. 941–958, Aug. 1989.
[7] A. R. Calderbank and L. H. Ozarow, "Non-equiprobable signaling on the Gaussian channel," *IEEE Trans. Inform Theory*, vol. 36, pp. 726–740, July 1990.
[8] G. D. Forney, Jr., "Trellis shaping," *IEEE Trans. Inform. Theory*, vol. 38, pp. 281–300, Mar. 1992.
[9] F. R. Kschischang and S. Pasupathy, "Optimal shaping properties of the truncated polydisc," *IEEE Trans. Inform. Theory*, Mar. 1992, submitted for publication.
[10] G. D. Forney, Jr., "Coset codes I: Introduction and geometrical classification," *IEEE Trans. Inform. Theory*, vol. 34, pp. 1123–1151, Sept. 1988.
[11] ——, "Coset codes II: Binary lattices and related codes," *IEEE Trans. Inform. Theory*, vol. 34, pp. 1152–1187, Sept. 1988.
[12] A. R. Calderbank and N. J. A. Sloane, "New trellis codes based on lattices and cosets," *IEEE Trans. Inform. Theory*, vol. IT-33, pp. 177–195, Mar. 1987.
[13] G. D. Forney, Jr., R. G. Gallager, G. R. Lang, F. M. Longstaff, and S. U. Qureshi, "Efficient modulation for band-limited channels," *IEEE J. Select. Areas Commun.*, vol. SAC-2, pp. 632–647, Sept. 1984.
[14] R. G. Gallager, "Source coded modulation system," U.S. Patent 4 586 182, Apr. 29, 1986.
[15] J. N. Livingston, "Shaping using variable-size regions," *IEEE Trans. Inform. Theory*, pp. 1347–1353, July 1992.
[16] A. Chouly and H. Sari, "Block-coded modulation: Novel design techniques and rotational invariance," *Philips J. Res.*, vol. 45, no. 2, pp. 127–155, 1990.
[17] G. R. Lang, G. D. Forney, S. Qureshi, F. M. Longstaff, and C. H. Lee, "Signal structures with data encoding/decoding for QCM modulations," U.S. Patent 4 538 284, Aug. 27, 1985.
[18] T. Armstrong, "Secondary channel signaling in a QAM data point constellation," U.S. Patent 4 630 287, Dec. 16, 1986.
[19] R. D. Gitlin and J.-J. Werner, "Inband coding of secondary data," U.S. Patent 4 644 537, Feb. 17, 1987.
[20] H. K. Thapar, "Inband coding of secondary data," U.S. Patent 4 651 320, Mar. 17, 1987.
[21] R. D. Gitlin, H. K. Thapar, and J. J. Werner, "An inband coding method for the transmission of secondary data," in *Conf. Rec. IEEE Int. Conf. Commun.*, June 1988, pp. 3.1.1–3.1.5.
[22] C. Hastings, Jr., *Approximations for Digital Computers.* Princeton, NJ: Princeton University Press, 1955.
[23] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions.* New York: Dover, 1965.
[24] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups.* New York: Springer-Verlag, 1988.
[25] L.-F. Wei, "Trellis-coded modulation with multidimensional constellations," *IEEE Trans. Inform. Theory*, vol. IT-33, pp. 483–501, July 1987.
[26] T. Berger, "Minimum entropy quantizers and permutation codes," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 149–157, Mar. 1982.
[27] I. Ingemarsson, "Optimized permutation modulation," *IEEE Trans. Inform. Theory*, pp. 1098–1100, Sept. 1990.
[28] J. A. Hertz, R. G. Palmer, and A. S. Krogh, *Introduction to the Theory of Neural Computation.* Reading, MA: Addison-Wesley, 1991.
[29] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671–680, May 13, 1983.
[30] N. J. A. Sloane, "Tables of sphere packings and spherical codes," *IEEE Trans. Inform. Theory*, vol. IT-27, pp. 327–338, May 1981.
[31] J. M. Wozencraft and I. M. Jacobs, *Principles of Communication Engineering.* New York: Wiley, 1965.
[32] W. Rudin, *Function Theory in Polydiscs.* New York: W. A. Benjamin, 1969.
[33] R. E. Blahut, *Digital Transmission of Information.* Reading, MA: Addison-Wesley, 1990.
[34] P. R. Stubley and I. F. Blake, "On a discrete probability distribution matching problem," *J. Algorithms*, June 1991, submitted for publication.
[35] D. A. Huffman, "A method for the construction of minimum redundancy codes," *Proc. IRE*, vol. 40, pp. 1098–1101, 1952.
[36] R. W. Hamming, *Coding and Information Theory.* Englewood Cliffs, NJ: Prentice-Hall, 1986.
[37] G. Ungerboeck, "Channel coding with multilevel/phase signals," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 55–67, Jan. 1982.